


Image Filtering & Edge Detection

Reading:
 □ Chapter 7 and 8, F&P

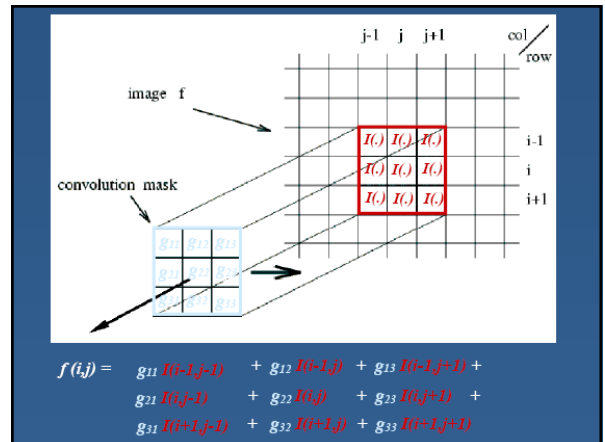
What is image filtering?

- Modify the pixels in an image based on some function of a local neighborhood of the pixels.



Linear Functions


- Simplest: linear filtering.
 - Replace each pixel by a linear combination of its neighbors.
- The prescription for the linear combination is called the “convolution kernel”.



Convolution

- Let I be the image and g be the kernel. The output of convolving I with g is denoted $I * g$

$$f[m,n] = I * g = \sum_{k,l} I[m-k,n-l]g[k,l]$$

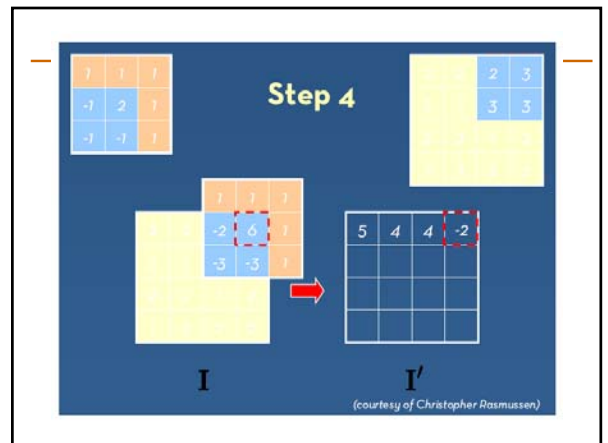
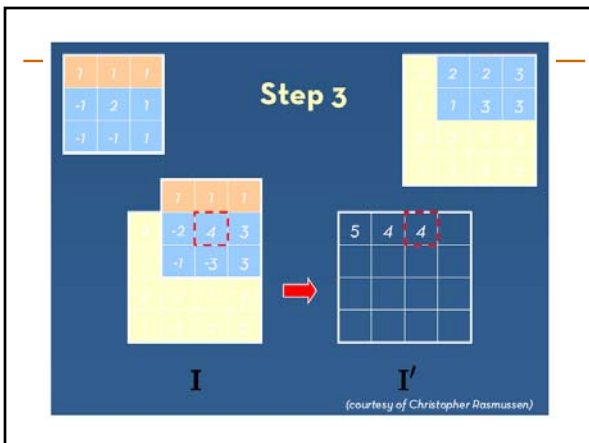


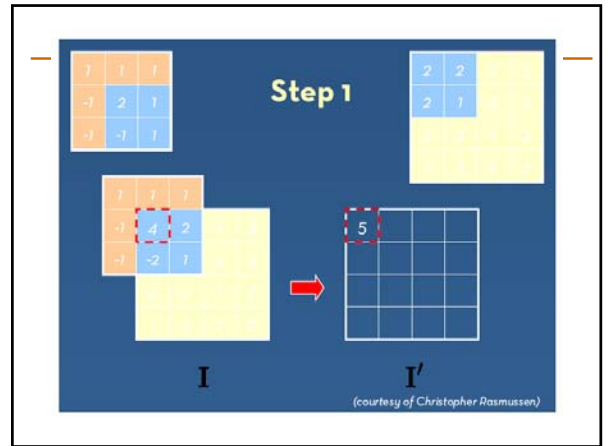


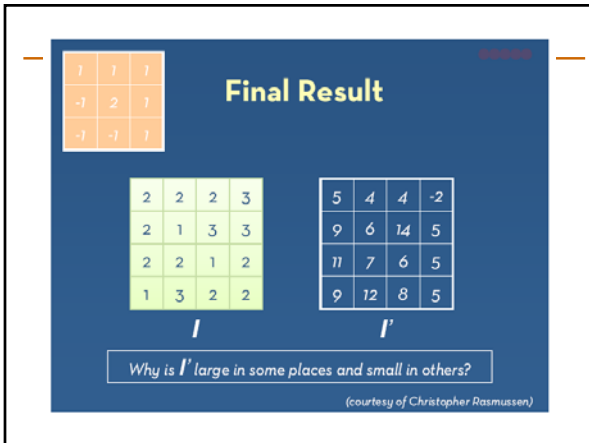

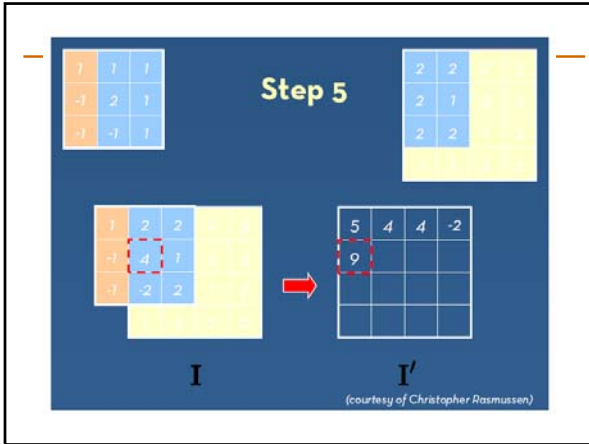
Key properties

- Linearity:**
 - $filter(I_1 + I_2) = filter(I_1) + filter(I_2)$
- Shift invariance:**
 - same behavior regardless of pixel location
 - $filter(shift(I)) = shift(filter(I))$
- Theoretical result:**
 - Any linear shift-invariant operator can be represented as a convolution**

Properties in more detail


- Commutative: $a * b = b * a$
 - Conceptually no difference between filter and signal
- Associative: $a * (b * c) = (a * b) * c$
 - Often apply several filters one after another: $((a * b_1) * b_2) * b_3$
 - This is equivalent to applying one filter: $a * (b_1 * b_2 * b_3)$
- Distributes over addition: $a * (b + c) = (a * b) + (a * c)$
- Scalars factor out: $ka * b = a * kb = k(a * b)$
- Identity: unit impulse $e = [\dots, 0, 0, 1, 0, 0, \dots]$,
 $a * e = a$





Yucky details

- What is the size of the output?
- MATLAB: `filter2(g, I, shape)`
 - `shape = 'full'`: output size is sum of sizes of I and g
 - `shape = 'same'`: output size is same as I
 - `shape = 'valid'`: output size is of difference sizes for I & g



Implementation details

- What about near the edge?
 - the filter window falls off the edge of the image
 - need to extrapolate
 - methods:
 - clip filter (black)
 - wrap around
 - copy edge
 - reflect across edge



Source: S. Marschner

Implementation details


- What about near the edge?
 - the filter window falls off the edge of the image
 - need to extrapolate
 - methods (MATLAB):
 - clip filter (black): `imfilter(f, g, 0)`
 - wrap around: `imfilter(f, g, 'circular')`
 - copy edge: `imfilter(f, g, 'replicate')`
 - reflect across edge: `imfilter(f, g, 'symmetric')`

Source: S. Marschner

Linear filtering (warm-up slide)



original




?

Linear filtering (warm-up slide)



original




Filtered
(no change)

| | | |
|---|---|---|
| 0 | 0 | 0 |
| 0 | 1 | 0 |
| 0 | 0 | 0 |

Linear filtering



original




?

Shift



original




shifted

| | | |
|---|---|---|
| 0 | 0 | 0 |
| 0 | 0 | 1 |
| 0 | 0 | 0 |

Linear filtering



original




?

Blurring



original




Blurred (filter applied in **both dimensions**).

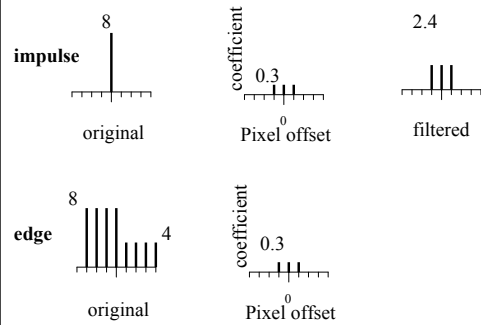
Box filter:

| | | |
|---|---|---|
| 1 | 1 | 1 |
| 1 | 1 | 1 |
| 1 | 1 | 1 |

Blur examples




Blur examples




Smoothing with box filter revisited

- Smoothing with an average actually doesn't compare at all well with a defocused lens
- Most obvious difference is that a single point of light viewed in a defocused lens looks like a fuzzy blob; but the averaging process would give a little square



Smoothing with box filter revisited


- Smoothing with an average actually doesn't compare at all well with a defocused lens
- Most obvious difference is that a single point of light viewed in a defocused lens looks like a fuzzy blob; but the averaging process would give a little square
- Better idea: to eliminate edge effects, weight contribution of neighborhood pixels according to their closeness to the center, like so:



"fuzzy blob"

Gaussian Kernel

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




- Constant factor at front makes volume sum to 1 (can be ignored, as we should re-normalize weights to sum to 1 in any case)

Source: C. Rasmussen

Choosing kernel width

- Gaussian filters have infinite support, but discrete filters use finite kernels



Source: K. Grauman

Gaussian filtering

- A Gaussian kernel gives less weight to pixels further from the center of the window


| | | | | | | | | | |
|---|---|---|----|----|----|----|----|---|---|
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | 90 | 90 | 90 | 90 | 90 | 0 | 0 |
| 0 | 0 | 0 | 90 | 90 | 90 | 90 | 90 | 0 | 0 |
| 0 | 0 | 0 | 90 | 90 | 90 | 90 | 90 | 0 | 0 |
| 0 | 0 | 0 | 90 | 90 | 90 | 90 | 90 | 0 | 0 |
| 0 | 0 | 0 | 90 | 90 | 90 | 90 | 90 | 0 | 0 |
| 0 | 0 | 0 | 90 | 90 | 90 | 90 | 90 | 0 | 0 |
| 0 | 0 | 0 | 90 | 90 | 90 | 90 | 90 | 0 | 0 |
| 0 | 0 | 0 | 90 | 90 | 90 | 90 | 90 | 0 | 0 |

$$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} H[u, v]$$


$F[x, y]$

- This kernel is an approximation of a Gaussian function:


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



Example: Smoothing with a Gaussian



Mean vs. Gaussian filtering



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

Source: D. Lowe

Separability example

2D convolution (center location only)

$$\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} * \begin{bmatrix} 2 & 3 & 3 \\ 3 & 5 & 5 \\ 4 & 4 & 6 \end{bmatrix}$$

The filter factors into a product of 1D filters:

$$\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} * \begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$$

Perform convolution along rows:

$$\begin{bmatrix} 1 & 2 & 1 \end{bmatrix} * \begin{bmatrix} 2 & 3 & 3 \\ 3 & 5 & 5 \\ 4 & 4 & 6 \end{bmatrix} = \begin{bmatrix} 11 & & \\ 18 & & \\ 18 & & \end{bmatrix}$$

Followed by convolution along the remaining column:

Source: K. Grauman

Gaussian filters


- Remove "high-frequency" components from the image (low-pass filter)
- Convolution with self is another Gaussian
 - So can smooth with small-width kernel, repeat, and get same result as larger-width kernel would have
 - Convoluting two times with Gaussian kernel of width σ is same as convoluting once with kernel of width $\sqrt{2}\sigma$
- Separable kernel
 - Factors into product of two 1D Gaussians
 - Useful: can convolve all rows, then all columns
 - How does this change the computational complexity?
 - Linear vs. quadratic in mask size

Source: K. Grauman

Review: Linear filtering

- What are the defining mathematical properties of a convolution?
- What is the difference between blurring with a box filter and blurring with a Gaussian?
- What happens when we convolve a Gaussian with another Gaussian?
- What is separability?
- How does separability affect computational complexity?

Noise




- Salt and pepper noise:** contains random occurrences of black and white pixels
- Impulse noise:** contains random occurrences of white pixels
- Gaussian noise:** variations in intensity drawn from a Gaussian normal distribution

Source: S. Seitz

Gaussian noise


- Mathematical model: sum of many independent factors
- Good for small standard deviations
- Assumption: independent, zero-mean noise



Local maxima Median operation Gaussian filter ("active") noise


Source: K. Grauman

Reducing Gaussian noise



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


Reducing salt-and-pepper noise



- 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?
 - No. \leftrightarrow Not a convolution

Source: K. Grauman

Median filter

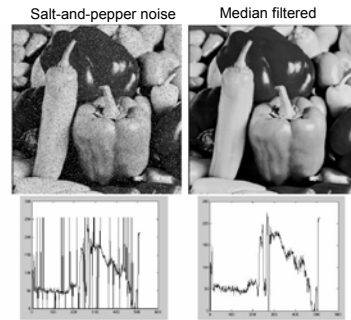
- What advantage does median filtering have over Gaussian filtering?
 - Robustness to outliers

filters have width 5 :



Source: K. Grauman

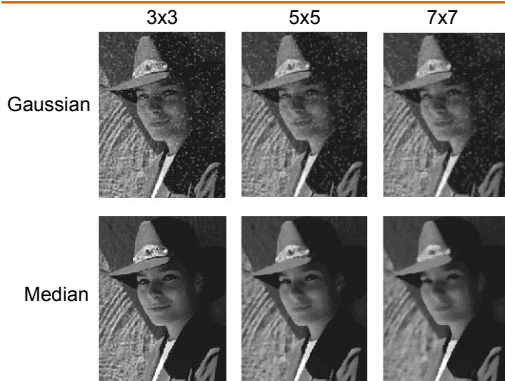
Median filter




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

Source: K. Grauman


Median vs. Gaussian filtering




Linear filtering (warm-up slide)



Linear filtering (no change)




Linear filtering



(Remember blurring)



original



Blurred (filter applied in both dimensions).

Sharpening



original




Sharpened original

Sharpening



original





Sharpened original

Sharpening example



original



Sharpened
(differences are accentuated; constant areas are left untouched).

Sharpening



Original

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

(Note that filter sums to 1)

?

Source: D. Lowe

Sharpening



Original


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



Sharpening filter
- Accentuates differences with local average

Source: D. Lowe

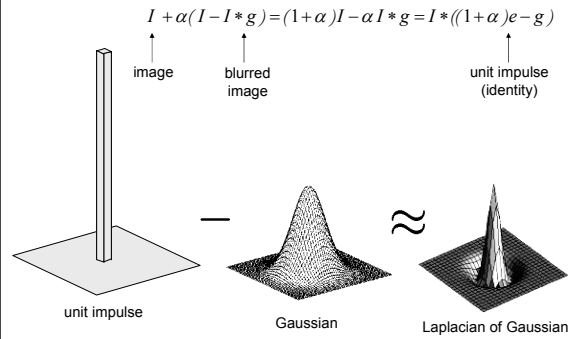
Sharpening



before

after

Unsharp mask filter



Sharpening Revisited

- What does blurring take away?



- Let's add it back:




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)



Source: D. Lowe

Origin of Edges



- Edges are caused by a variety of factors

Source: Steve Seitz

Characterizing edges

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





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



- The gradient direction is given by:

$$\theta = \tan^{-1} \left(\frac{\frac{\partial f}{\partial y}}{\frac{\partial f}{\partial x}} \right)$$
 - how does this relate to the direction of the edge? *perpendicular*
- The *edge strength* is given by the gradient magnitude

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

Source: D. Forsyth, D. Lowe

Differentiation and convolution

- Recall, for 2D function, $f(x,y)$:
- We could approximate this as


$$\frac{\partial f}{\partial x} = \lim_{\epsilon \rightarrow 0} \left(\frac{f(x+\epsilon, y) - f(x, y)}{\epsilon} \right) \quad \frac{\partial f}{\partial x} \approx \frac{f(x_{n+1}, y) - f(x_n, y)}{\Delta x}$$

- This is linear and shift invariant, so must be the result of a convolution.
- (which is obviously a convolution)



Source: D. Forsyth, D. Lowe

Finite differences: example



- Which one is the gradient in the x-direction (resp. y-direction)?

Finite difference filters

- Other approximations of derivative filters exist:

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



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

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

Source: K. Grauman

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?

Effects of noise

- Finite difference filters respond strongly to noise
 - Image noise results in pixels that look very different from their neighbors
 - Generally, the larger the noise the stronger the response
- What is to be done?

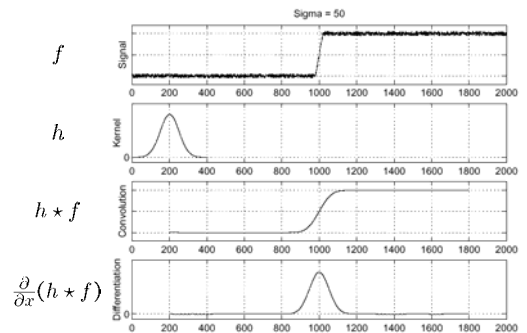
Source: D. Forsyth

Effects of noise

- Finite difference filters respond strongly to noise
 - Image noise results in pixels that look very different from their neighbors
 - Generally, the larger the noise the stronger the response
- What is to be done?
 - Smoothing the image should help, by forcing pixels different to their neighbors (=noise pixels?) to look more like neighbors

Source: D. Forsyth

Solution: smooth first




- Where is the edge? ■ Look for peaks in $\frac{\partial}{\partial x}(h * f)$

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:




Source: S. Seitz

Derivative of Gaussian filter



Derivative of Gaussian filter




- Which one finds horizontal/vertical edges?

Summary: Filter mask properties

- Filters act as templates
 - Highest response for regions that “look the most like the filter”
 - Dot product as correlation
- Smoothing masks
 - Values positive
 - Sum to 1 → constant regions are unchanged
 - Amount of smoothing proportional to mask size
- Derivative masks
 - 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

Source: K. Grauman

Tradeoff between smoothing and localization



1 pixel

3 pixels

7 pixels

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

Source: D. Forsyth

Implementation issues




- The gradient magnitude is large along a thick “trail” or “ridge,” so how do we identify the actual edge points?
- How do we link the edge points to form curves?

Source: D. Forsyth


Laplacian of Gaussian

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



- Where is the edge? Zero-crossings of bottom graph

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

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

Laplacian of Gaussian

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


∇^2 is the Laplacian operator:

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

MATLAB demo

```
g = fspecial('gaussian',15,2);
imagesc(g)
surf1(g)
gclown = conv2(clown,g,'same');
imagesc(conv2(clown,[-1 1],'same'));
imagesc(conv2(gclown,[-1 1],'same'));
dx = conv2(g,[-1 1],'same');
imagesc(conv2(clown,dx,'same'));
lg = fspecial('log',15,2);
lclown = conv2(clown,lg,'same');
imagesc(lclown)
imagesc(clown + .2*lclown)
```


Edge finding




We wish to mark points along the curve where the magnitude is biggest. We can do this by looking for a maximum along a slice normal to the curve (non-maximum suppression). These points should form a curve. There are then two algorithmic issues: at which point is the maximum, and where is the next one?

Source: D. Forsyth

Non-maximum suppression




At q , we have a maximum if the value is larger than those at both p and at r . Interpolate to get these values.




Source: D. Forsyth

Predicting the next edge point




Assume the marked point is an edge point. Then we construct the tangent to the edge curve (which is normal to the gradient at that point) and use this to predict the next points (here either r or s).



Source: D. Forsyth

Designing an edge detector

- Criteria for an "optimal" edge detector:
 - **Good detection:** the optimal detector must minimize the probability of false positives (detecting spurious edges caused by noise), as well as that of false negatives (missing real edges)
 - **Good localization:** the edges detected must be as close as possible to the true edges
 - **Single point:** the detector must minimize the number of responses per edge



Source: L. Fei-Fei

Canny edge detector

- This is probably the most widely used edge detector in computer vision
- Theoretical model: step-edges corrupted by additive Gaussian noise
- Canny has shown that the first derivative of the Gaussian closely approximates the operator that optimizes the product of *signal-to-noise ratio* and localization

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

Source: L. Fei-Fei


Canny edge detector

1. Filter image with derivative of Gaussian
2. Find magnitude and orientation of gradient
3. Non-maximum suppression:
 - Thin multi-pixel 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')`

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



- thinning
■ (non-maximum suppression)

Hysteresis thresholding



original image



high threshold
(strong edges)



low threshold
(weak edges)



hysteresis threshold

Source: L. Fei-Fei

Effect of σ (Gaussian kernel spread/size)

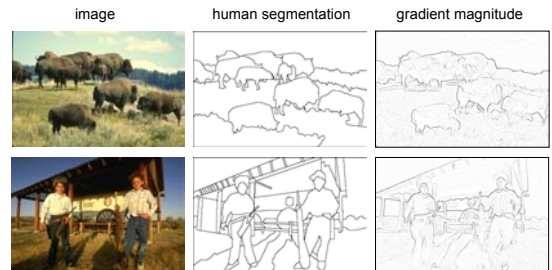


original Canny with $\sigma = 1$ Canny with $\sigma = 2$

- The choice of σ depends on desired behavior
 - large σ detects large scale edges
 - small σ detects fine features

Source: S. Seitz

Edge detection is just the beginning...



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