Image segmentation

GW Chapter 10

Pratt Chapter 7

1

Image segmentation

- Goal: partition the image into its constituent objects
- Approaches
 - Discontinuity: detect abrupt changes in gray levels \rightarrow edge detection
 - Similarity: group pixels based on their similarity with respect to a predefined criterion → region-based processing
 - Feature extraction
 - Region growing
 - Feature clustering/classification



Edge detection

Digital Image Processing K. Pratt, Chapter 15, pag 443 GW Chapter 10

Edge detection

- Framework: image segmentation
- Goal: identify objects in images
 - but also feature extraction, multiscale analysis, 3D reconstruction, motion recognition, image restoration, registration
- Classical definition of the edge detection problem: localization of large local changes in the grey level image → large graylevel gradients
 - This definition does not apply to *apparent edges*, which require a more complex definition
 - Extension to color images
- Contours are very important perceptual cues!
 - They provide a first *saliency map* for the interpretation of image semantics



Contours as perceptual cues



What do we detect?

- Depending on the impulse response of the filter, we can detect different types of graylevel discontinuities
 - Isolate points (pixels)
 - Lines with a predefined slope
 - Generic contours
- However, edge detection implies the evaluation of the local gradient and corresponds to a (directional) derivative

Detection of Discontinuities

• Point Detection



Detection of Discontinuities

• Line detection



Detection of Discontinuities

Line Detection Example: •

bc



Threshold=max{filtered} Suitable for binary images





Continuous domain edge models





2D discrete domain single pixel spot models

а а aaaa а aaaa а а a a a baaa a a а а a a а а a aaaaaa aaaaaa Step spot а a aaaaa a a a a а a a aacccaa a a cbcaa а acccaa aaaaa а а aaaaaa Single pixel transition spot a a a a a a а а a a a a d d a a d a a a a а d aaaaa a a aaaaaa Smoothed transition spot







Models of an ideal digital edge



Types of edge detectors

- Unsupervised or autonomous: only rely on local image features
 - No contextual information is accounted for
 - Simple to implement, flexible, suitable for generic applications
 - Not robust
- Supervised or contextual: exploit other sources of information
 - Some a-priori knowledge on the semantics of the scene
 - Output of other kind of processing
 - Less flexible
 - More robust
- There is no golden rule: the choice of the edge detection strategy depends on the application

Types of edge detection

- Differential detection
 - Differential operators are applied to the original image F(x,y) to produce a differential image G(x,y) with accentuated spatial amplitude changes
 - Thresholds are applied to select locations of large amplitude
- Model fitting
 - Fitting a local region of pixel values to a model of the edge, line or spot
- A binary indicator map E(x,y) is used to indicate the positions of edges, lines or points



• First order derivatives

• Second order derivatives



Diff. edge det.: Approach

- 1. Smoothing of the image
 - To reduce the impact of noise and the number of spurious (non meaningful) edges
 - To regularize the differentiation
- 2. Calculation of first and second order derivatives
 - Isolation of high spatial frequencies
 - Required features: invariance to rotations, linearity
 - Critical point: choice of the scale (size of the support)
- 3. Labeling
 - Plausibility measure for the detected point belonging to a contour (to get rid of false edges)

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

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

• The edge strength is given by the gradient magnitude

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



Orthogonal gradient vector

• Continuous 1D gradient along a line normal to the edge slope

$$G(x, y) = \frac{\partial f}{\partial x} \cos \theta + \frac{\partial f}{\partial y} \sin \theta$$

• Need of a discrete approximation: definition of a row and a column gradient combined in a spatial gradient amplitude

$$G[j,k] = \left(\left| G_{row}[j,k] \right|^2 + \left| G_{col}[j,k] \right|^2 \right)^{1/2}$$

$$G[j,k] = \left| G_{row}[j,k] \right| + \left| G_{col}[j,k] \right|$$

$$\Im[j,k] = \arctan\left\{ \frac{G_{col}[j,k]}{G_{row}[j,k]} \right\}$$
computationally more efficient

orientation of the spatial gradient with respect to the row axis



Simplest row/col gradient approximations k k $G_{row}[j,k] \cong f[j,k] - f[j,k-1]$ -1 1 $G_{col}[j,k] \cong f[j,k] - f[j+1,k]$ -1 k vertical step edge model: a a a a b b b b b 0000h0000 vertical ramp edge model: a a a a c b b b b 0 0 0 0 (h/2 h/2) 0 0 0 c=(a+b)/2→ the edge is not located $G_{row}[j,k] \cong f[j,k+1] - f[j,k-1]$ at the midpoint of the ramp $G_{col}[j,k] \cong f[j-1,k] - f[j+1,k]$ 0 0 h/2 h h/2 0 0

The discrete gradient

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

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

- Discrete approximation

0	-1	0	-1	-1	-1
-1	4	-1	-1	8	-1
0	-1	0	-1	-1	-1

Example



(a) Original



FIGURE 15.2-2. Horizontal and vertical differencing gradients of the peppers_mon image.

Diagonal gradients

• Robert's cross-difference operator

$$G[j,k] = \left(\left| G_R[j,k] \right|^2 + \left| G_C[j,k] \right|^2 \right)^{1/2}$$
$$G[j,k] = \left| G_R[j,k] \right| + \left| G_C[j,k] \right|$$
$$\vartheta[j,k] = \frac{\pi}{4} \arctan\left\{ \frac{G_C[j,k]}{G_R[j,k]} \right\}$$

square root form

magnitude form

edge orientation with respect to the row axis

 $G_{R}[j,k] = f[j,k] - f[j+1,k+1]$ $G_{C}[j,k] = f[j,k+1] - f[j+1,k]$



Example: Robert's



(a) Magnitude

(b) Square root

FIGURE 15.2-3. Roberts gradients of the peppers_mon image.

Orthogonal differential gradient edge op.

	o]	0	٢٥	٥٦	-1	٥٦
Pixel difference	0	1	-1	0	1	0
	Lo	0	ο]	Lo	0	0
	۲°	0	٥٦	o ٦	-1	٢٥
Separated	1	0	-1	0	0	0
pixel difference	_ o	0	ο]	Lo	1	0
	۲°	0	-1]	[-1	0	٢٥
Roberts	0	1	0	0	1	0
	Lo	0	۵]	Lo	0	٥J
	.[1	0	-1]	.[^{−1}	-1	-1]
Prewitt	$\frac{1}{2}$ 1	0	-1	$\frac{1}{2}$ 0	0	0
	3 1	0	_1]	°[1	1	1
	¹].	0	-1]	.[^{−1}	-2	-1]
Sobel	1 2	0	-2	- 0	0	0
	*L 1	0	_1]	"L 1	2	1
	, [¹	0	-1]	· [−1	$-\sqrt{2}$	-17
Frei-Chen	$\frac{1}{\sqrt{2}}$	0	$-\sqrt{2}$	- 0	0	0
	$2 + \sqrt{2} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$	0	_1]	$2 + \sqrt{2} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$	$\sqrt{2}$	1

Gradients as convolutions

 The gradient calculation is a neighborhood operation, so it can be put in matrix notations

 $G_{row}[j,k] = f[j,k] * H_{row}[j,k]$ $G_{col}[j,k] = f[j,k] * H_{col}[j,k]$

- H_{row/col}: row and column impulse response arrays

- The size of the convolution kernels can be increased to improve robustness to noise
- Example: normalized boxcar operator

$$\mathbf{H}_{R} = \frac{1}{21} \begin{bmatrix} 1 & 1 & 1 & 0 & -1 & -1 & -1 \\ 1 & 1 & 1 & 0 & -1 & -1 & -1 \\ 1 & 1 & 1 & 0 & -1 & -1 & -1 \\ 1 & 1 & 1 & 0 & -1 & -1 & -1 \\ 1 & 1 & 1 & 0 & -1 & -1 & -1 \\ 1 & 1 & 1 & 0 & -1 & -1 & -1 \\ 1 & 1 & 1 & 0 & -1 & -1 & -1 \end{bmatrix}$$



Introducing averaging

 Differential methods are highly sensitive to small luminance fluctuations → combine with averaging



• Prewitt operator square root gradient

$$\begin{aligned} \mathsf{k=1} & G(j,k) = \left[\left[G_R(j,k) \right]^2 + \left[G_C(j,k) \right]^2 \right]^{1/2} \\ G_R(j,k) &= \frac{1}{K+2} \left[(A_2 + KA_3 + A_4) - (A_0 + KA_7 + A_6) \right] \\ G_C(j,k) &= \frac{1}{K+2} \left[(A_0 + KA_1 + A_2) - (A_6 + KA_5 + A_4) \right] \end{aligned}$$

Sobel, Frei&Chen operator

- Sobel: same as Prewitt with k=2
 - Give the same importance to each pixel in terms to its contribution to spatial gradient
- Frei&Chen: same as Prewitt with k=sqrt(2)
 - The gradient is the same for horizontal, vertical and diagonal edges
Sobel



a ₀	a 1	a ₂
a ₇	(i,j)	a ₃
a ₆	a ₅	a_4

Special case of the general one hereafter with c=2

$$G_{row} [i,j] = (a0 + c \ a7 + a6) - (a2 + c \ a3 + a4)$$

$$G_{col} = (a6 + c \ a5 + a4) - (a0 + c \ a1 + a2)$$

$$c=2$$

$$G = \sqrt{G_{row}^2 + G_{col}^2}$$





- Kirsch operator
 - 8 directional masks, each selecting one specific direction
 - "winner takes all" paradigm for the absolute value of the gradient and direction selected by the index of the corresponding mask
- Robinson operator
 - 8 directional masks, similar to Kirsh





Truncated pyramid op.

Linearly decreasing weighting to pixels away from the center of the edge

$$\mathbf{H}_{R} = \frac{1}{34} \begin{bmatrix} 1 & 1 & 1 & 0 & -1 & -1 & -1 \\ 1 & 2 & 2 & 0 & -2 & -2 & -1 \\ 1 & 2 & 3 & 0 & -3 & -2 & -1 \\ 1 & 2 & 3 & 0 & -3 & -2 & -1 \\ 1 & 2 & 3 & 0 & -3 & -2 & -1 \\ 1 & 2 & 2 & 0 & -2 & -2 & -1 \\ 1 & 1 & 1 & 0 & -1 & -1 & -1 \end{bmatrix}$$



Edge detection in presence of noise



FIGURE 10.7 First column: images and gray-level profiles of a ramp edge corrupted by a random Gaussian noise of mean 0 and $\sigma = 0.0, 0.1, 1.0, and 10.0, respectively. Second col$ b umn: first-derivative images and gray-level profiles. Third column: second-derivative с images and gray-level profiles.

d

Improving robustness to noise

- Combining smoothing with differentiation
 - Solution 1: do smoothing first and then differentiation
 - Solution 2: differentiate the smoothing filter and do filtering



h: smoothing filter

Solution 1: Smoothing+Differentiation



Sol. 2: Differentiation of the smoothing filter





47

Extending to 2° order derivative

• The derivative of a convolution is equal to the convolution of either of the functions with the derivative of the other

$$z(x) = f(x) * g(x)$$
$$\frac{dz}{dx} = \frac{df}{dx} * g = f * \frac{dg}{dx}$$

• Iterating

$$z(x) = f(x) * g(x)$$
$$\frac{d^2 z}{dx^2} = \frac{d}{dx} \left(\frac{df}{dx} * g\right) = \frac{d^2 f}{dx^2} * g$$

Hints of the proof

• Intuition (OP)

$$\begin{aligned} c(t) &= f(t)^* g(t) = f^* g(t) = \int_{-\infty}^{+\infty} f(\tau) g(t-\tau) d\tau \\ c'(t) &= \frac{dc(t)}{dt} = \frac{d}{dt} \Big(f(t)^* g(t) \Big) \\ C(\omega) &= \Im \{ c(t) \} = \Im \{ f(t)^* g(t) \} = F(\omega) G(\omega) \\ \Im \{ c'(t) \} &= j \omega \Im \{ c(t) \} = j \omega F(\omega) G(\omega) = \begin{cases} [j \omega F(\omega)] G(\omega) \to f'(t)^* g(t) \\ F(\omega) [j \omega G(\omega)] \to f(t)^* g'(t) \end{cases} \end{aligned}$$

Remark

- The order in which differentiation and smoothing are performed depends on their properties.
 - Such operations are interchangeable as long as they are linear. Thus, if both smoothing and differentiation are performed by linear operators they are interchangeable
 - In this case they can be performed at the same time by filtering the image with the differentiation of the smoothing filter

- Argyle and Macleod
- Laplacian of Gaussian (LoG)
 - Difference of Gaussians (DoG)

Argyle and Macleod

Use a large neighborhood Gaussian-shaped weighting for noise suppression

$$g(x,s) = [2\pi s^2]^{-1/2} \exp\{-1/2(x/s)^2\}$$

Argyle operator horizontal coordinate impulse response array can be expressed as a sampled version of the continuous domain impulse response. s and t are the spread parameters

Argyle operator
horizontal coordinate
impulse responseArgyle
 $H_R(j,k) = \begin{cases} -2g(x,s)g(y,t) & \text{for } x \ge 0 \\ \\ 2g(x,s)g(y,t) & \text{for } x < 0 \end{cases}$ for $x \ge 0$ array $H_R(j,k) = \begin{cases} -2g(x,s)g(y,t) & \text{for } x < 0 \\ \\ 2g(x,s)g(y,t) & \text{for } x < 0 \end{cases}$ smoothing along the
edge and differentiation
along the orthogonal

direction

McLeod $H_R(j,k) = [g(x+s,s)-g(x-s,s)]g(y,t)$

The Argyle and Macleod operators, unlike the boxcar operator, give decreasing importance to pixels far removed from the center of the neighborhood.

Argyle and Macleod

- Extended-size differential gradient operators can be considered to be compound operators in which a smoothing operation is performed on a noisy image followed by a differentiation operation.
- The compound gradient impulse response can be written as

 $H(j,k) \,=\, H_G(j,k) \,\circledast\, H_S(j,k)$

gradient op. low pass

• Example

if Hg is the 3x3 Prewitt row gradient operator and Hs (j,k) = 1/9, for all (j,k) in a 3x3 matrix, is a uniform smoothing operator, the resultant row gradient operator, after normalization to unit positive and negative gain, becomes

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$$
$$\mathbf{H}_{R} = \frac{1}{18} \begin{bmatrix} 1 & 1 & 0 & -1 & -1 \\ 2 & 2 & 0 & -2 & -2 \\ 3 & 3 & 0 & -3 & -3 \\ 2 & 2 & 0 & -2 & -2 \\ 1 & 1 & 0 & -1 & -1 \end{bmatrix}$$

Second order derivative

- Edge detectors based on first order derivative are not robust
 - High sensitivity to noise, need a threshold
- Second order derivative operators detect the edge at the zerocrossing of the second derivative → more robust, more precise
 - Less sensitive to noise, usually don't need a threshold for postprocessing of the contours image



Laplace operator

• Second order differentiation operator

$$\Delta f = \nabla^2 f = \nabla \cdot \nabla f$$
$$\nabla^2 f = \sum_{i=1}^N \frac{\partial^2 f}{\partial x_i^2}$$
$$N = 2 \longrightarrow \nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

• Directional derivative

$$D_{\vec{v}}f(\vec{x}) = \sum_{i=1}^{N} v_i \frac{\partial f}{\partial x_i}$$
$$v_i = \left\langle \vec{v}, \vec{i} \right\rangle$$

Laplace operator

• Second order derivative in the continuous domain

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

• Discrete approximation

$$\frac{\partial^2 f}{\partial x^2} = \frac{\partial G_x}{\partial x} = \frac{\partial}{\partial x} \left[f(i, j+1) - f(i, j) \right] =$$

$$= \frac{\partial f(i, j+1)}{\partial x} - \frac{\partial f(i, j)}{\partial x} =$$

$$= \left[f(i, j+2) - f(i, j+1) \right] - \left[f(i, j+1) - f(i, j) \right] =$$

$$= f(i, j+2) - 2f(i, j+1) + f(i, j)$$

Discrete approximation: proof

• Centring the estimation on *(i,j)* the simplest approximation is to compute the difference of slopes along each axis

 $G(x, y) = -\nabla^2 f(x, y)$

$$\begin{split} G_{row}[i,j] &= \left(f[i,j] - f[i,j-1]\right) - \left(f[i,j+1] - f[i,j]\right) = 2f[i,j] - f[i,j-1] - f[i,j+1] \\ G_{col}[i,j] &= \left(f[i,j] - f[i+1,j]\right) - \left(f[i-1,j] - f[i,j]\right) = 2f[i,j] - f[i+1,j] - f[i-1,j] \end{split}$$

• This can be put in operator and matrix form as

$$G[i, j] = f[i, j] * H[i, j]$$
$$H = \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$$

Discrete approximation

- The 4-neighbors Laplacian is often normalized to provide unit gain averages of the positive and negative weighted pixels in the 3x3 neighborhood
- Gain normalized 4-neighbors Laplacian

$$H = \frac{1}{4} \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$$

- The weights of the pixels in the neighborhood, and thus the normalization coefficient, can be changed to emphasize the edges.
 - Ex. Prewitt modified Laplacian

$$H = \frac{1}{8} \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

Discrete approximation

- Gain normalized separable 8 neighbors Laplacian

$$H = \frac{1}{8} \begin{bmatrix} -2 & 1 & -2 \\ 1 & 4 & 1 \\ -2 & 1 & -2 \end{bmatrix}$$

а	а	а	b	b
а	а	а	b	b
а	а	а	b	b
а	а	а	b	b
а	а	а	b	b

Note

- Without sign change after the evaluation of the Laplacian
 - However, the sign is meaningless if we evaluate the modulus of the gradient

$$\frac{\partial^2 f}{\partial x^2} = f(i, j+1) - 2f(i, j) + f(i, j-1)$$
$$\frac{\partial^2 f}{\partial y^2} = f(i+1, j) - 2f(i, j) + f(i-1, j)$$

• Different possible Laplacian matrices



Laplacian of Gaussian

- Quite often the zero crossing does not happen at a pixel location
 - See the example of the step edge
- It is common choice to locate the edge at a pixel with a positive response having a neighbor with a negative response
- Laplacian of Gaussian: Marr&Hildreth have proposed an operator in which Gaussian shaped smoothing is performed *prior* to the application of the Laplacian
- Continuous LoG gradient

$$LOG(x, y) = -\nabla^{2} \left\{ f(x, y) * H_{s}(x, y) \right\}$$
$$H_{s}(x, y) = g(x, s)g(y, s)$$
$$g(x, s) = \frac{1}{\sqrt{2\pi s^{2}}} \exp \left\{ -\frac{1}{2} \left(\frac{x}{s}\right)^{2} \right\} \qquad \text{in } G_{k}(x, y) = \frac{1}{\sqrt{2\pi s^{2}}} \exp \left\{ -\frac{1}{2} \left(\frac{x}{s}\right)^{2} \right\}$$

impulse response of the Gaussian smoothing kernel

LoG operator

 As a result of the linearity of the second derivative operator and of the convolution

$$LOG[j,k] = f[j,k] * H[j,k]$$
(1)
$$H(x,y) = -\nabla^{2} \left\{ g(x,s)g(y,s) \right\}$$
$$H(x,y) = \frac{1}{\pi s^{4}} \left(1 - \frac{x^{2} + y^{2}}{2s^{2}} \right) \exp \left\{ -\frac{x^{2} + y^{2}}{2s^{2}} \right\}$$

- It can be shown that
 - The convolution (1) can be performed separately along rows and cols
 - It is possible to approximate the LOG impulse response closely by a difference of Gaussians (DOG) operator

 $H(x, y) = g(x, s_1)g(y, s_1) - g(x, s_2)g(y, s_2), \quad s_1 < s_2$

The LoG operator

$$g(x, y) = \frac{1}{2\pi s^2} \exp\left[-\frac{x^2 + y^2}{2s^2}\right]$$

$$h(x, y) : \nabla^2 \left[g(x, y) * f(x, y)\right] = \left[\nabla^2 g(x, y)\right] * f(x, y) = h(x, y) * f(x, y)$$

where

$$\nabla^2 g(x, y) = \frac{x^2 + y^2 - 2s^2}{2\pi s^4} \exp\left[-\frac{x^2 + y^2}{2s^2}\right] \text{ mexican hav}$$

- How to choose s?
 - Large values: pronounced smoothing → better denoising BUT smears out sharp boundaries reducing the precision in edge localization
 - Small values: soft smoothing \rightarrow lower noise reduction BUT better boundary preservation
 - A good solution could be to follow a multiscale approach (s is the scale)

LoG filtering

- Gaussian smoothing (low-pass filter)
 - Noise reduction (the larger the filter, the higher the smoothing)
 - BUT
 - Smears out edges
 - Blurs the image (defocusing)
- Laplacian detection (high-pass filter)
- Edge location by interpolation
 - The zero-crossing does not happen in a pixel site

LoG filtering = Gaussian smoothing + Laplacian detection



FDoG

- First derivative of Gaussian op. [Pratt]
 - Gaussian shaped smoothing is followed by differentiation
 - FDoG continuous domain horizontal impulse response

$$H_{R}(j,k) = \frac{-\partial [g(x,s)g(y,t)]}{\partial x}$$

$$g(x,s) = [2\pi s^{2}]^{-1/2} \exp\{-1/2(x/s)^{2}\}$$

$$H_{R}(j,k) = \frac{-xg(x,s)g(y,t)}{s^{2}} \qquad \text{introduces the sign change}$$

$$H_{R}(x,y) = -\frac{x}{s^{2}} g(x,s)g(y,t)$$

$$H_{C}(x,y) = -\frac{y}{t^{2}} g(x,s)g(y,t)$$





LoG

- Independent variables
 - s value: larger values allow larger denoising but smear out details and made contour extraction not quite precise
- Solutions
 - Trade off
 - Multiscale

2D edge detection filters



LoG: example

• The Laplacian of a Gaussian filter



A digital approximation:

0	0	1	0	0
0	1	2	1	0
1	2	-16	2	1
0	1	2	1	0
0	0	1	0	0

Second derivative

 Laplacian of Gaussian: (LoG) – Mexican Hat

0	1	0
1	-4	1
0	1	0

 Laplacian of Gaussian: Link to early vision: the 2D Mexican Hat closely resembles the receptive fields of simple cells in the retina
 → edge detection is one of the first steps in vision



Laplacian zero-crossing detection

- Zero-valued Laplacian response pixels are unlikely in real images
- Practical solution: form the maximum of all positive Laplacian responses and the minimum of all Laplacian responses in a 3x3 window. If the *difference* between the two exceeds a threshold an edge is assumed to be present.
- Laplacian zero-crossing patterns



+: zero or positive
Laplacian of Gaussian (LoG)



Zero-crossings of bottom graph

Effects of noise

- Consider a single row or column of the image
 - Plotting intensity as a function of position gives a signal



Gradient thresholding



Revisiting Line detection

 Possible filters to find gradients along vertical and horizontal directions



Edge Detection

a b c d FIGURE 10.10 (a) Original image. (b) $|G_x|$, component of the gradient in the *x*-direction. (c) $|G_y|$, component in the *y*-direction. (d) Gradient image, $|G_x| + |G_y|$.



Edge Detection



FIGURE 10.11 Same sequence as in Fig. 10.10, but with the original image smoothed with a 5×5 averaging filter.

Edge Detection

a b

Laplacian mask



c d e f g

FIGURE 10.15 (a) Original image. (b) Sobel gradient (shown for comparison). (c) Spatial Gaussian smoothing function. (d) Laplacian mask. (e) LoG. (f) Thresholded LoG. (g) Zero crossings. (Original image courtesy of Dr. David R. Pickens, Department of Radiology and Radiological Sciences, Vanderbilt University Medical Center.)





One simple method to find zerocrossings is black/white thresholding:

- 1. Set all positive values to white
- 2. Set all negative values to black

3. Determine the black/white transitions.

Compare (b) and (g):

- Edges in the zero-crossings image is thinner than the gradient edges.
- Edges determined by zero-crossings have formed many closed loops.

Edge detection: Gradient thresholding

Prewitt filter: decreasing the threshold





Edge detection: Gradient thresholding

Prewitt filter: decreasing the threshold







Application to image enhancement







ſf (a) Input image; (b) Laplacian of (a); (c) Spatially invariant high-pass filtering [sum of (a) and (b)]; (d) Mask image [Sobel gradient of (a) smoothed by a 5x5 box filter]; (e) Product of (b) and (d); (f) Spacevariant enhancement [sum of (a) and (e)].

Multiscale edge detection

- The information obtained by filtering the image at different scales is combined to determine the edge map
 - scale \leftrightarrow width (s, sigma parameter) of the filter
- Different possibilities
 - Adapting the filter bandwidth to the local characteristics of the image (Wiener)
 - Combining edge maps obtained at different scales
- Canny algorithm
 - Smoothing (allows for different scales)
 - Gradient maxima
 - Two thresholds to detect both *weak* and *strong* edges. Weak edges are retained if they are connected to strong ones (labeling)
 - Less sensible to noise

Canny algorithm

- Based on a 1D continuous model of a step edge of amplitude h_E plus additive Gaussian noise of amplitude σ_n
- The impulse response of the filter h(x) is assumed to be FIR and *antisymmetric*
- *First order* derivative: the edge is located at the local maxima of

f(x) * h(x)

- A threshold has to be chosen
- Criterion: the Canny operator impulse response h(x) is chosen to satisfy three criteria
 - Good detection
 - Good localization
 - Single response

Step edge model

- Parameters
 - Edge direction (tangent to the curve)
 - Normal direction (vector orthogonal to the contour at edge location)
 - Local contrast (edge strength)
 - Edge location (along the normal direction)



Detection

- Criterion: The amplitude of the Signal to Noise Ratio (SNR) of the gradient is maximized for good detection
 - to obtain low probability of failure to mark edge points (*false negative rate*) and low probability to mark non-edge points (*false positive rate*)



Localization

- Criterion: Edge points marked by the ed operator must be as close as possible to the center of the edge
- Localization factor

$$LOC = \frac{h_E L(h)}{\sigma_n}$$
$$L(h) = \frac{h'(0)}{\int_{-W}^{W} [h'(x)]^2 dx}$$
$$h'(x) = \frac{dh(x)}{dx}$$

Single response

- Criterion: There should be only a single response to a true edge
 - The distance between peaks of the gradient when only noise is present is set to

$$x_m = kW \tag{2}$$

- Global criterion: maximization of the product S(h)L(h) subject to (2)
 - Constrained maximization
 - Note: a large filter (W) improves detection (better denoising) BUT reduces the precision in localization
 - No close form solution, numerical ones are adopted
 - For low x_m , h(x) resembles the boxcar, while for larger x_m it is closely approximated by a FDoG (first derivative of Gaussian)

Canny impulse response



FIGURE 15.2-8. Comparison of Canny and first derivative of Gaussian impulse response functions.





Performance assessment

- Possible errors
 - False negatives (an edge point is present but it is not detected)
 - False positives (a non-edge point is detected)
 - Error in the estimation of the orientation
 - Error in the localization of the edge
- Paradigms
 - Use of synthetic images + noise with known parameters
 - Tests on sets of real images

Performance evaluation

Objective

- The ground truth is assumed to be available and represented by the actual contour (*full reference metric*)
- Concerns low level features
 - Measure to which extent the estimated contour represents the actual contour
- Metric: MSE among the estimated (*f*[*j*,*k*]) and the real (*s*[*j*,*k*]) edges

Subjective

- The ground truth is not necessarily given (*reduced or no-reference metric*)
- Concerns high-level features
 - Measures to which extent the estimated contour allows to identify the corresponding object in the image
 - Focus on semantics or image content
- Metric: subjective scores given to the different algorithms
- Lead to perception-based models and metrics



- A common strategy in signal detection theory is to establish **a bound on the probability of false detection** resulting from noise and then try to maximize the probability of true signal detection
- When applied to edge detection, this translates in **setting a the** *minimum value of the threshold* such that the *FP rate does not exceed the predefined bound*. Then the probability of true edge detection can be calculated by a coincidence comparison of the edge maps of the ideal versus the real edge detectors

Performance assessment: Figure of Merit

- Types of errors
- Detection
 - Missing valid edge points (False Negatives, FN)
 - Failure to localize edge points
 - Classification of noise fluctuations as edge points (False Positives, FP)
- Localization
 - Error in estimating the edge angle;
 - Mean square distance of the edge estimate from the true edge
- Accuracy
 - Algorithm's tolerance to distorted edges and other features such as corners and junctions

Performance assessment: Figure of Merit



 I_{i} , I_{A} : number of ideal and detected edge points, respectively d_{i} : distance among the ideal and the detected edge point along the normal to a line of ideal edge points (evaluated according to (3)) α : scaling constant The rating factor is normalized so that R = 1 for a perfectly detected edge



Filters competition

- A possible classification strategy
 - Synthetic image
 - 64x64 pixels
 - vertical oriented edge with variable slope and contrast
 - added Gaussian noise of variance σ_n
 - Control parameter SNR=(h/ σ_n), h being the normalized edge value (0<h<=1)
 - Filter threshold: maximize the FM constrained to maximum bound for false detection rate
 - False detection=false positives
 - Probability to detect an edge when no edge is present

along a line



FIGURE 15.5-11. Edge location figure of merit for a vertical ramp edge as a function of signal-to-noise ratio for h = 0.1 and w = 1.



FIGURE 15.5-12. Edge location figure of merit for a vertical ramp edge as a function of signal-to-noise ratio for h = 0.1 and SNR = 100.

Changing SNR

- Sobel
- Step edge





(b) Edge map, R = 100%

(d) Edge map, R = 85.1%



SNR = 10 (c) Original



FIGURE 15.5-13. Edge location performance of Sobel edge detector as a function of signalto-noise ratio, h = 0.1, w = 1, a = 1/9.

Changing the filter





(b) East compass, R = 66.1%







(d) Roberts square root, R = 37.0%



(e) Sobel, R = 85.1%



FIGURE 15.5-14. Edge location performance of several edge detectors for SNR = 10, 103 h = 0.1, w = 1, a = 1/9.

Subjective evaluation



- Task: "Give a score to the detected edges"
- Many trials
 - The experiment is repeated at least two times for each subject
- Many subjects
 - A sufficiently high number of subjects must participate in the experiment to make data analysis significant from the statistical point of view
- Output: {scores}
- Data analysis

A high figure of merit generally corresponds to a well-located edge upon visual scrutiny, and vice versa.