Image Analysis and Processing – IM5-1IA

Segmentation - Basics

Laurent Najman

laurent.najman@esiee.fr

ESIEE Paris

Université Paris-Est, Laboratoire d'Informatique Gaspard-Monge, Équipe A3SI

Before we start

- How are we doing?
- Questions?
- Concerns?
- Tutorial problems?
- Project problems?

Syllabus

1. Introduction, image perception and representation
2. Enhancements – Histogram & pixelwise transforms
3. Transforms – FFT
4. Filtering – Linear filters
5. **Segmentation - Basics**
6. ?

Outline

1. Introduction
   - What is segmentation?
   - The two main approaches to segmentation
   - What do we want to achieve with segmentation
2. Pixel-based segmentation
3. Region-based segmentation

I – Introduction to segmentation

What is segmentation?

- Image segmentation is all about cutting the image into smaller *regions*.
- These new regions are supposed to define outlines of objects in the image.
- More often, segmented regions are just area of images with similar properties.
- Proper object recognition usually requires high-level knowlegde. This knowlegde can be built into the segmentation method, or used after the segmentation is done on the obtained regions.

How to achieve this ?

Tons of methods: Here are listed a few:

- Thresholding
- Histogram clustering
- Spatial clustering
- Split & Merge
- Active contours
- Level sets
- Knowledge based methods
- Morphological methods, among which region-growing methods
- Optimal methods

Top-down vs. bottom-up

- Top-down: apply segmentation method as an image simplification method, then sort out the mess.
- Bottom-up: carefully select constraints for segmentation based on a-priori knowledge, apply method, segmentation will be mostly right.
- Examples:
  - By design, snakes have to be used in BU fashion
  - Big selling point of level sets methods and optimal methods: can be used TD with similar regularity as the snake methods
  - Optimal methods integrate knowledge and other type of contraints
  - *Learning* is going into the way
  - MM methods regularize by constraints, not by method.

The two main approaches to segmentation

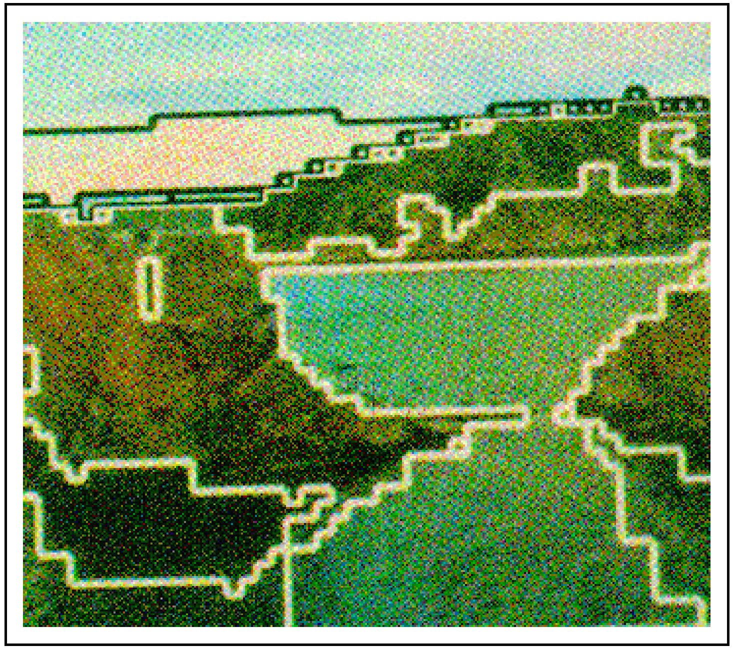

Top–down

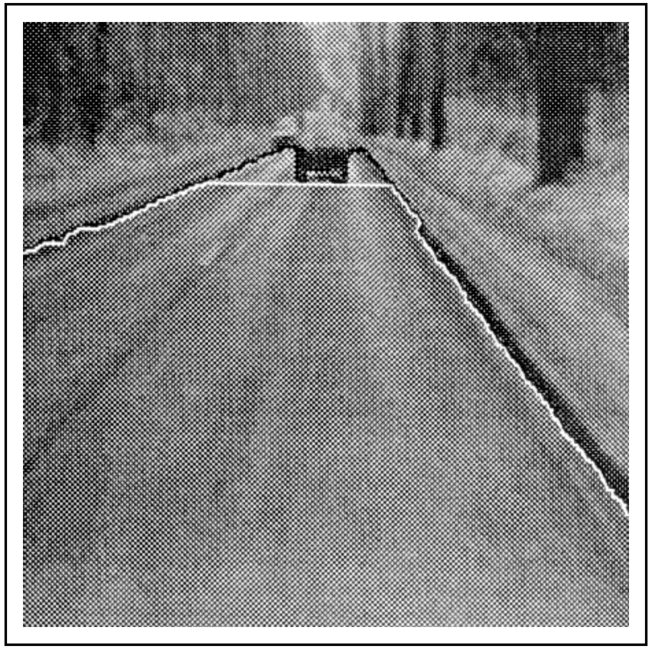

Bottom–up

Either top–down or bottom–up

II: Pixel-based segmentation

The simplest segmentation method?

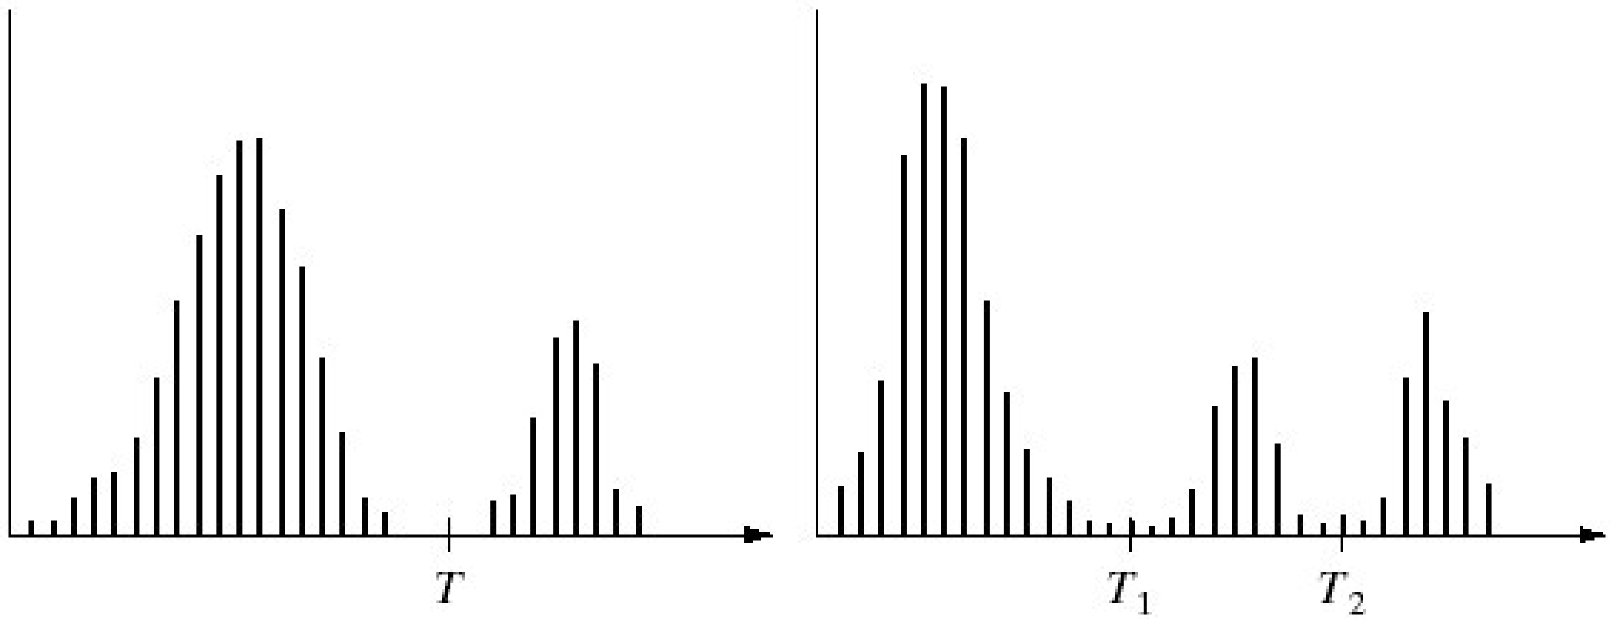

a b

FIGURE 10.26 (a) Gray-level histograms that can be partitioned by (a) a single threshold, and (b) multiple thresholds.

Thresholding model

FIGURE 10.32 Gray-level probability density functions of two regions in an image.

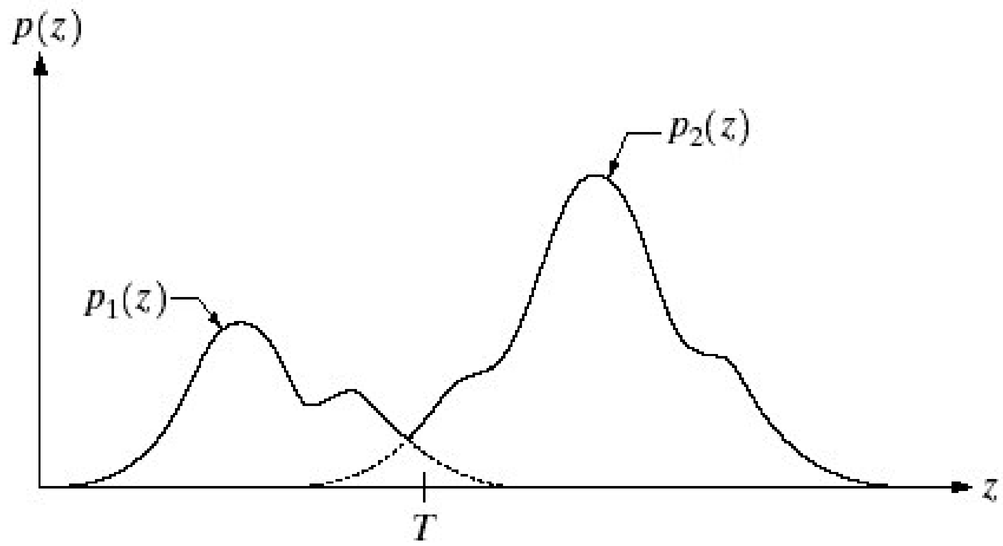

Thresholding model (cont.)

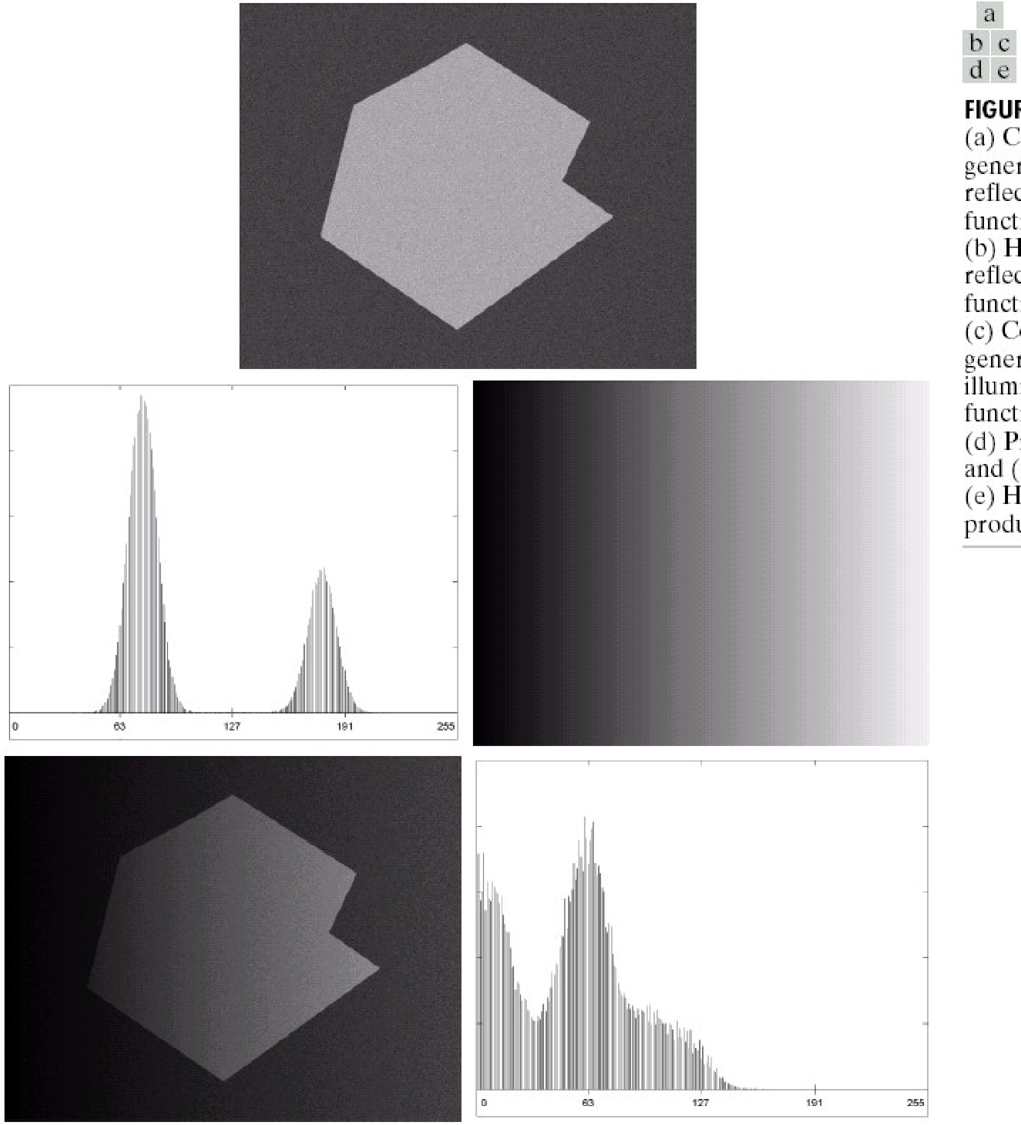

FIGURE 10.27
(a) Computer generated reflectance function.
(b) Histogram of reflectance function.
(c) Computer generated illumination function.
(d) Product of (a) and (c).
(e) Histogram of product image.

Object thresholding

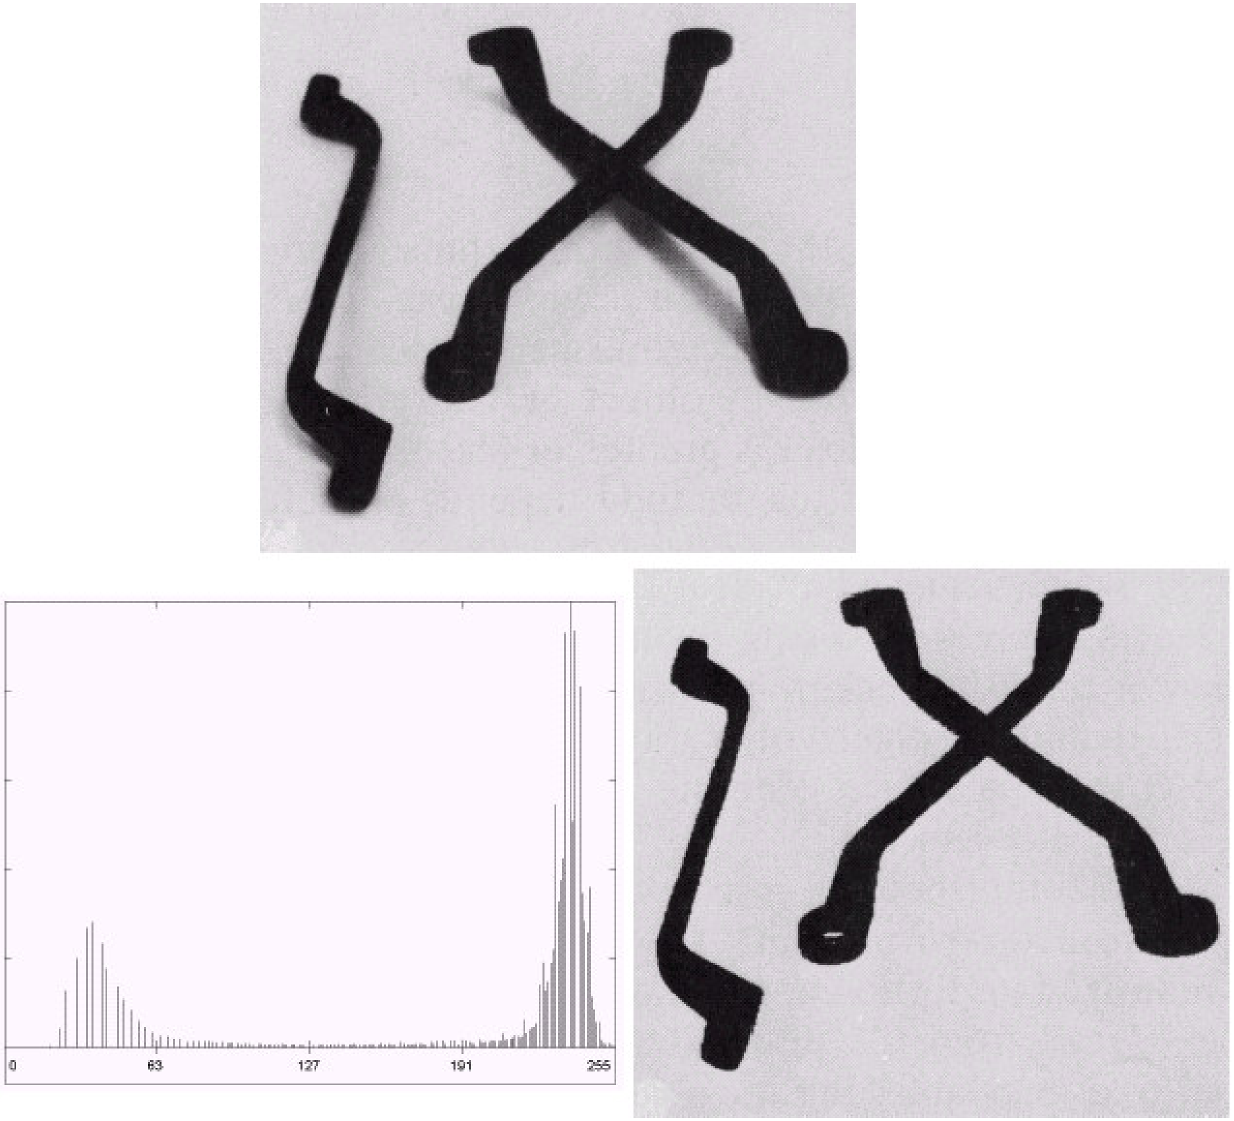

a
b c

FIGURE 10.28 (a) Original image. (b) Image histogram. (c) Result of global thresholding with T midway between the maximum and minimum gray levels.

Thresholding example

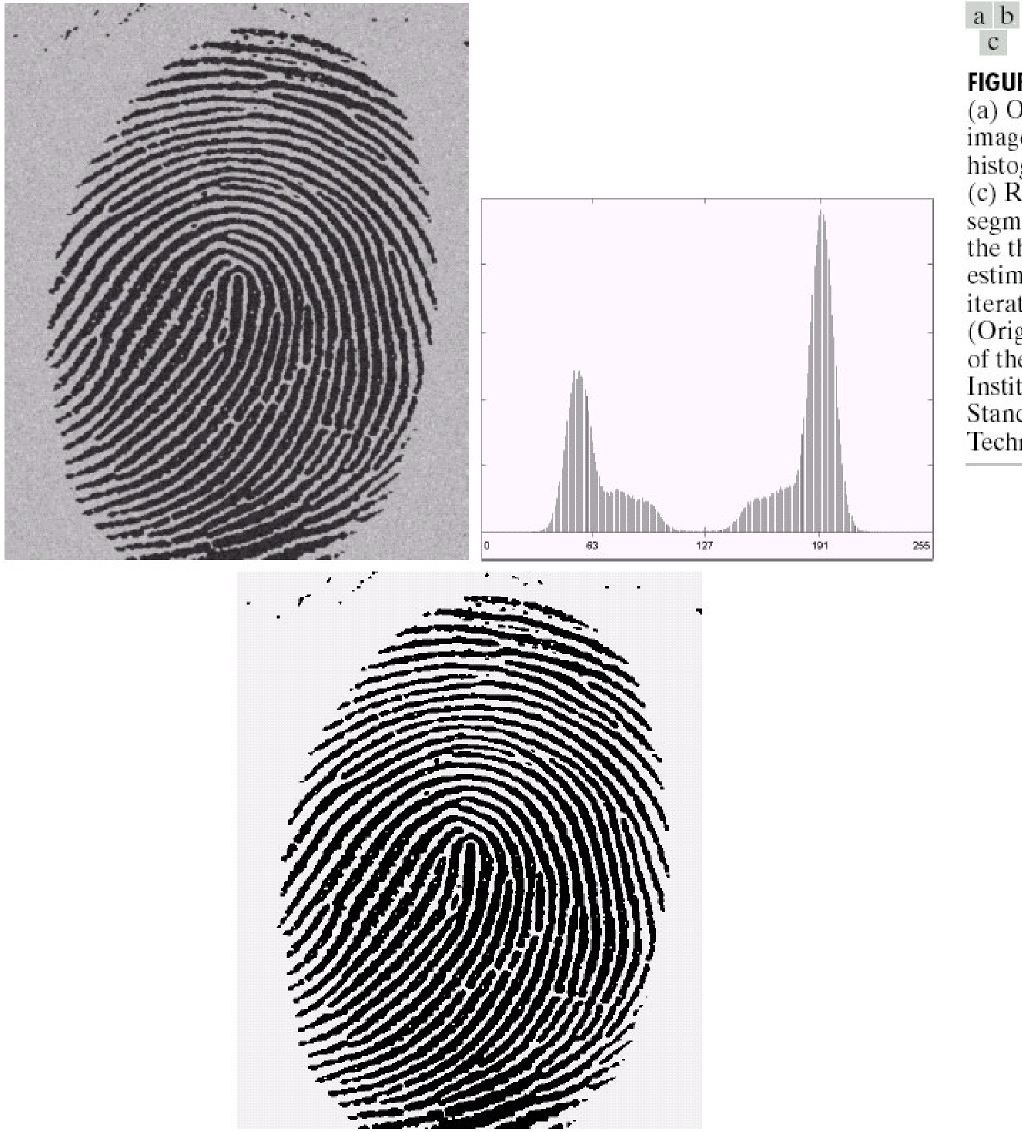

FIGURE 10.29 (a) Original image. (b) Image histogram. (c) Result of segmentation with the threshold estimated by iteration. (Original courtesy of the National Institute of Standards and Technology.)

Adaptive thresholding

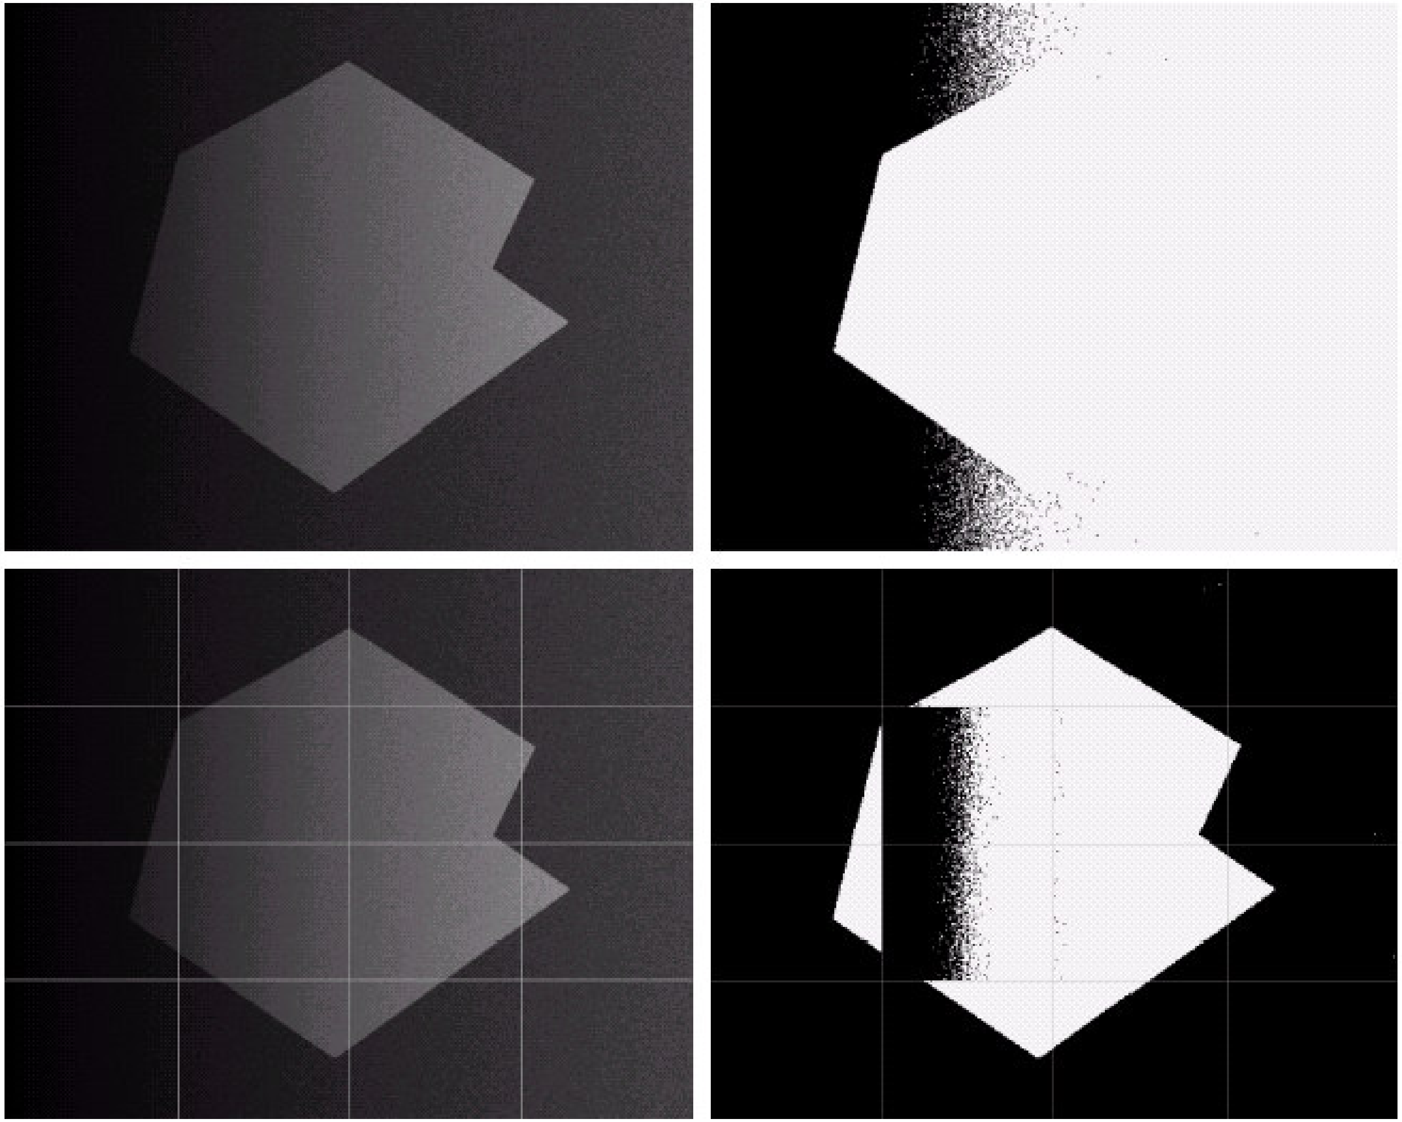

FIGURE 10.30 (a) Original image. (b) Result of global thresholding. (c) Image subdivided into individual subimages. (d) Result of adaptive thresholding.

Adaptive thresholding 2

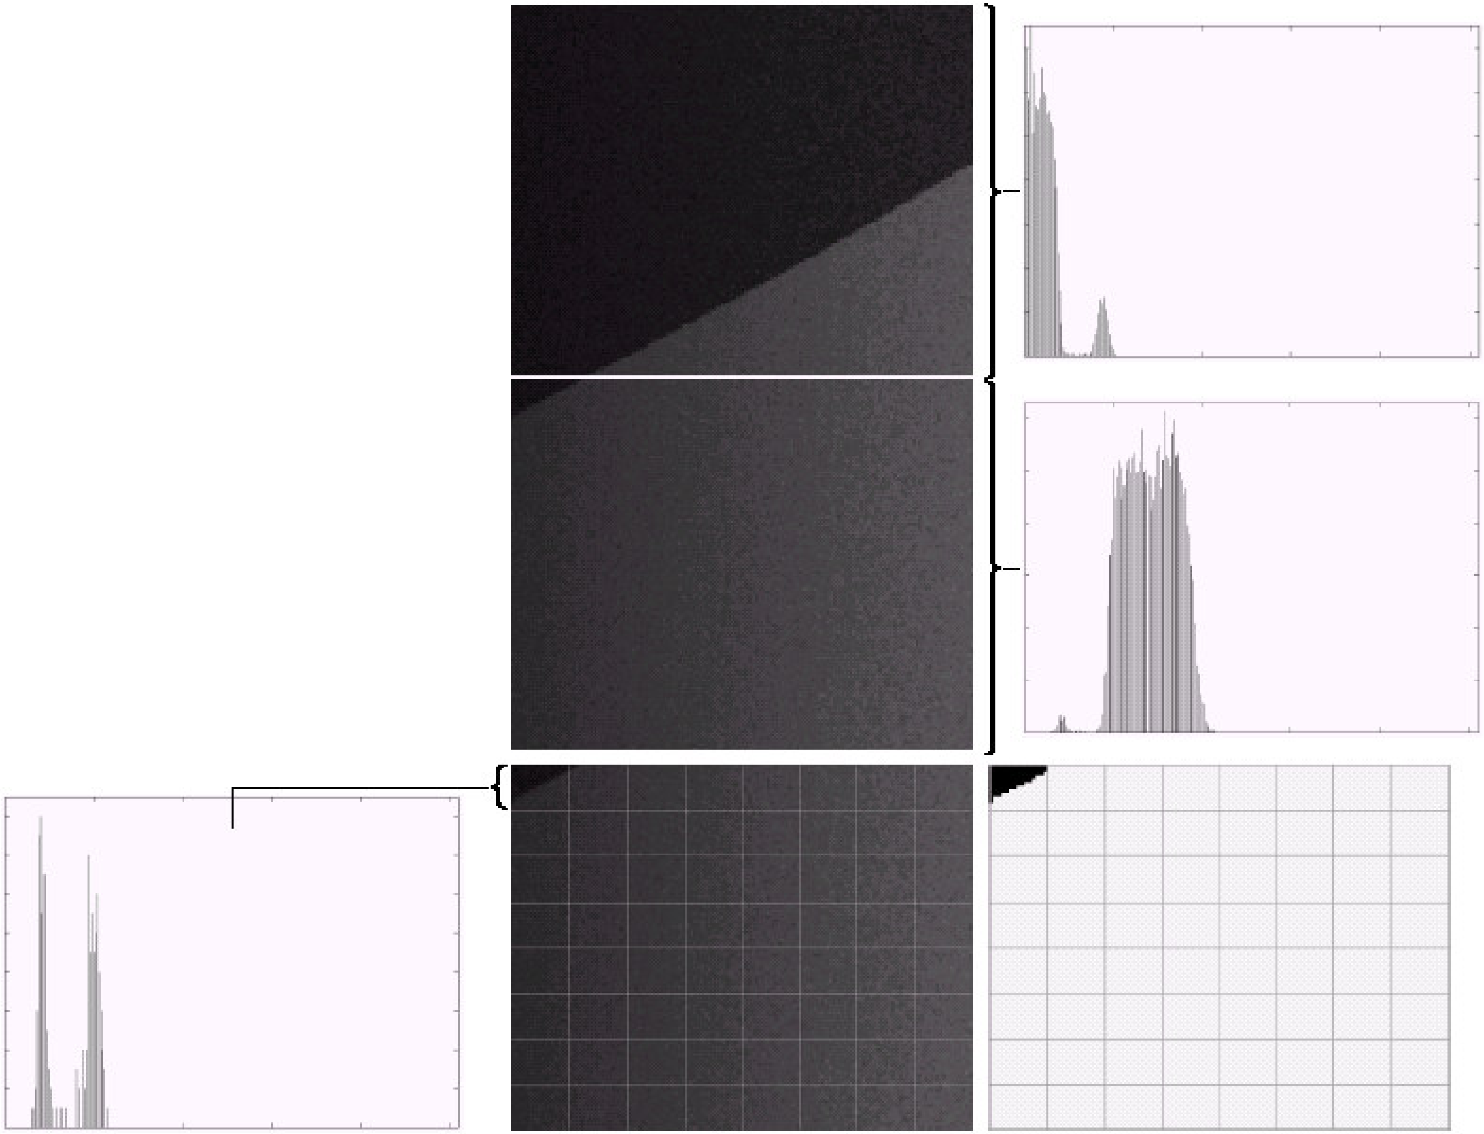

Adaptive thresholding 3

a
b

FIGURE 10.37
(a) Original image. (b) Image segmented by local thresholding. (Courtesy of IBM Corporation.)

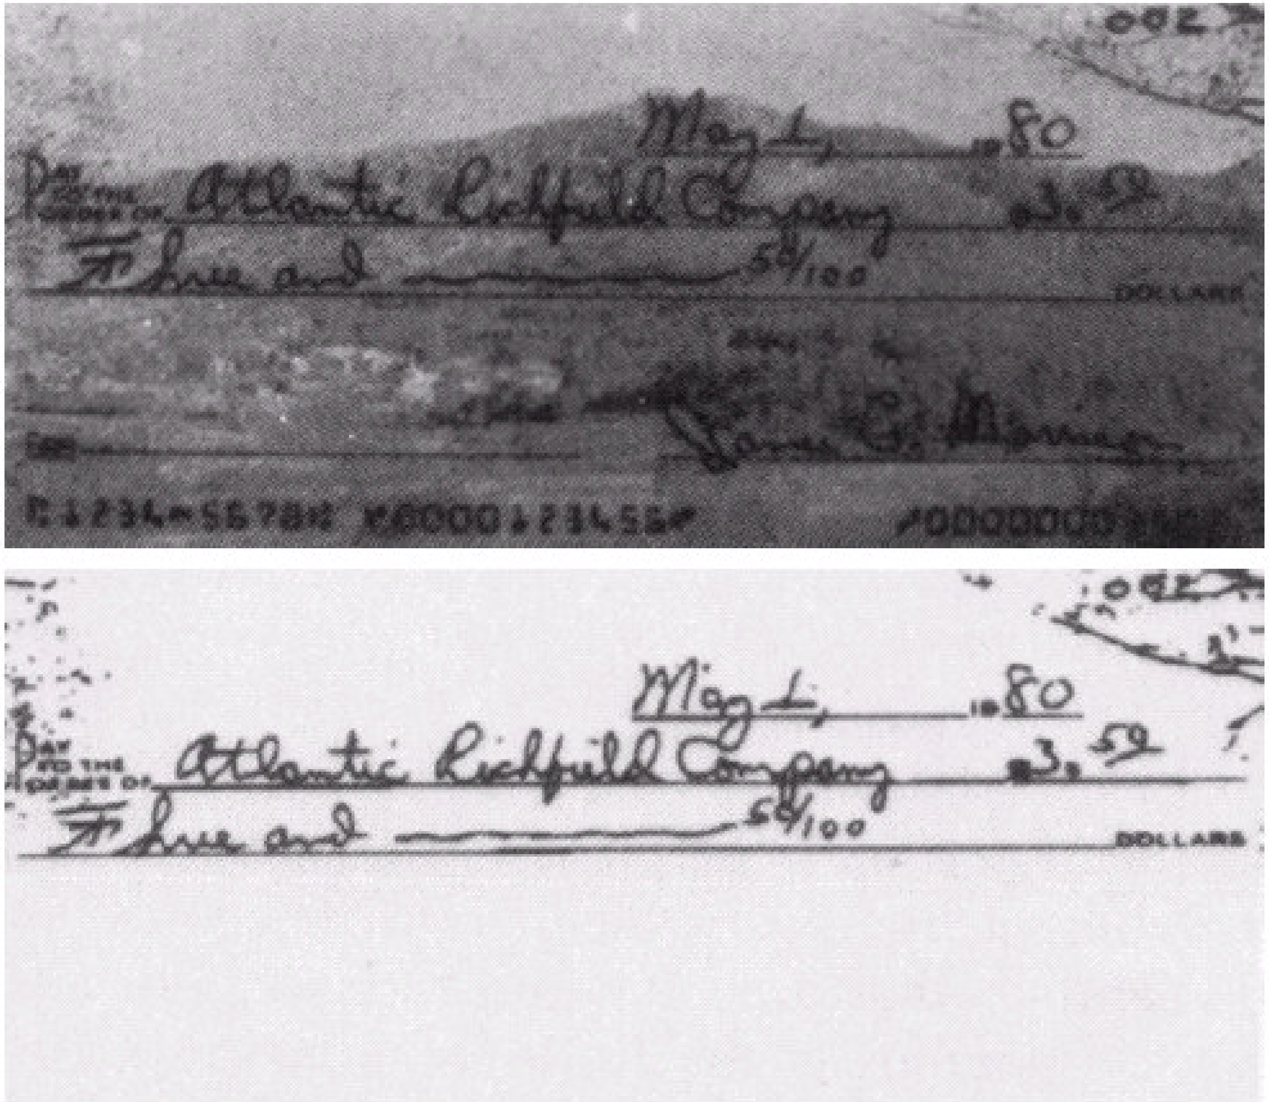

Multispectral thresholding

a b c

FIGURE 10.39 (a) Original color image shown as a monochrome picture. (b) Segmentation of pixels with colors close to facial tones. (c) Segmentation of red components.

Another example: skin lesions (1)

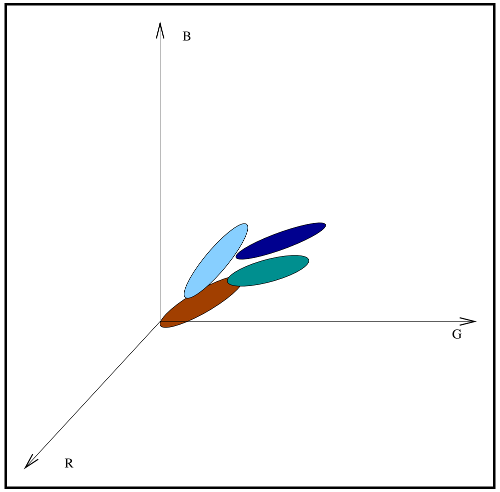

RGB space

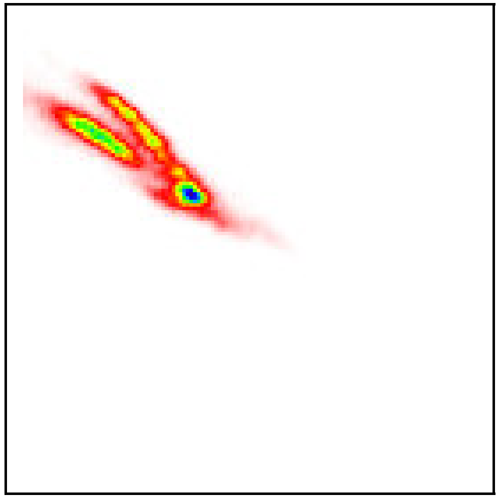

slice of RGB histo.

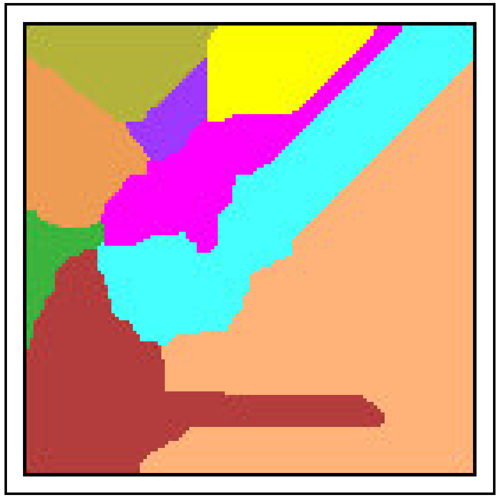

classification

Another example: skin lesions (2)

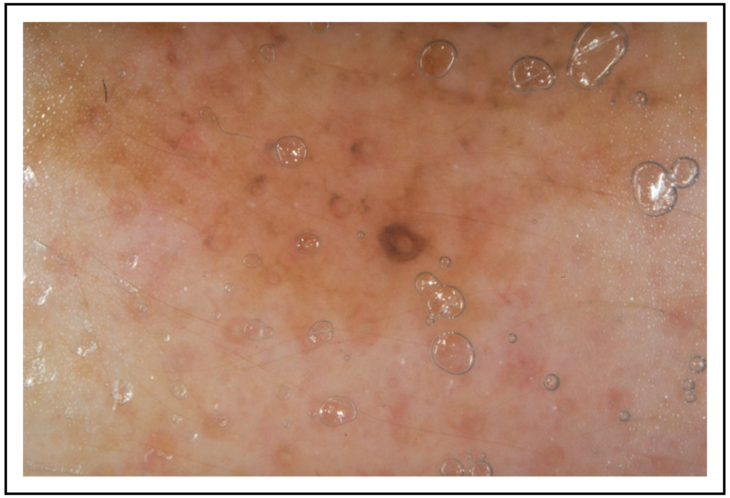

Original

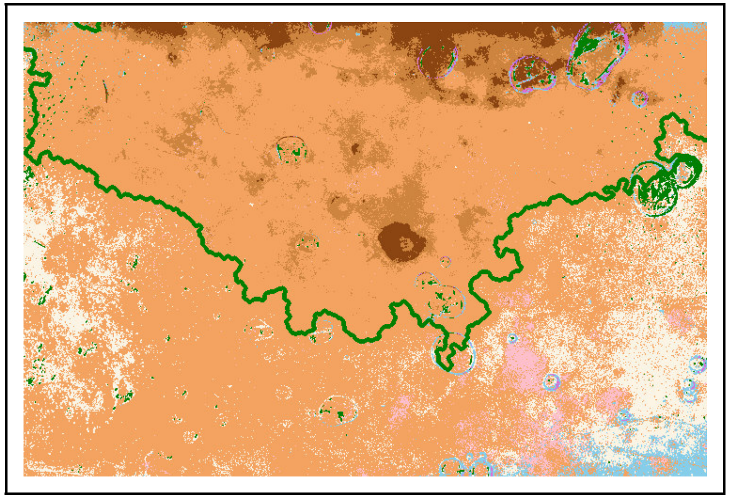

Colour segmentation

Summary on thresholding

- Many segmentation method end up using some thresholding
- Consequently many thresholding methods reduce to pre-processing followed by thresholding
- Adaptive thresholding model often too simplistic.

III: edge detection

Convolution-based methods

FIGURE 10.3 Line masks.

Horizontal

−1	−1	−1
2	2	2
−1	−1	−1

+45°

−1	−1	2
−1	2	−1
2	−1	−1

Vertical

−1	2	−1
−1	2	−1
−1	2	−1

−45°

2	−1	−1
−1	2	−1
−1	−1	2

Convolution-based example

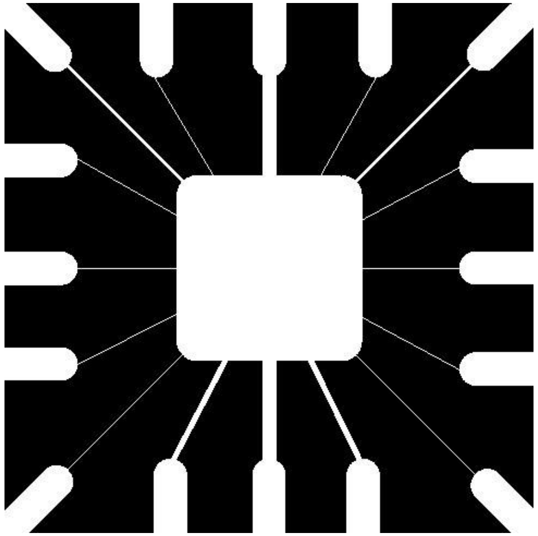

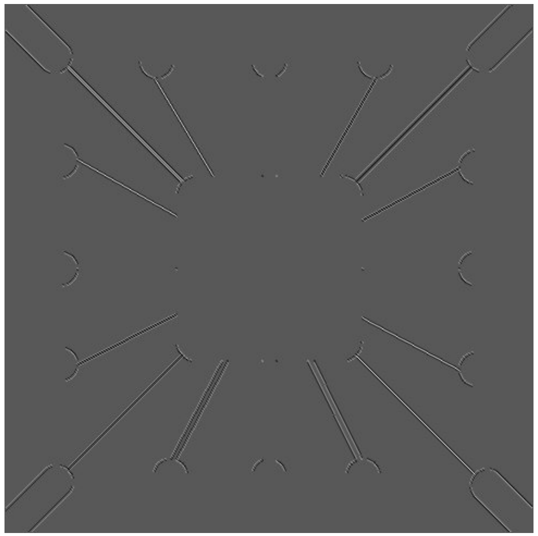

Line model

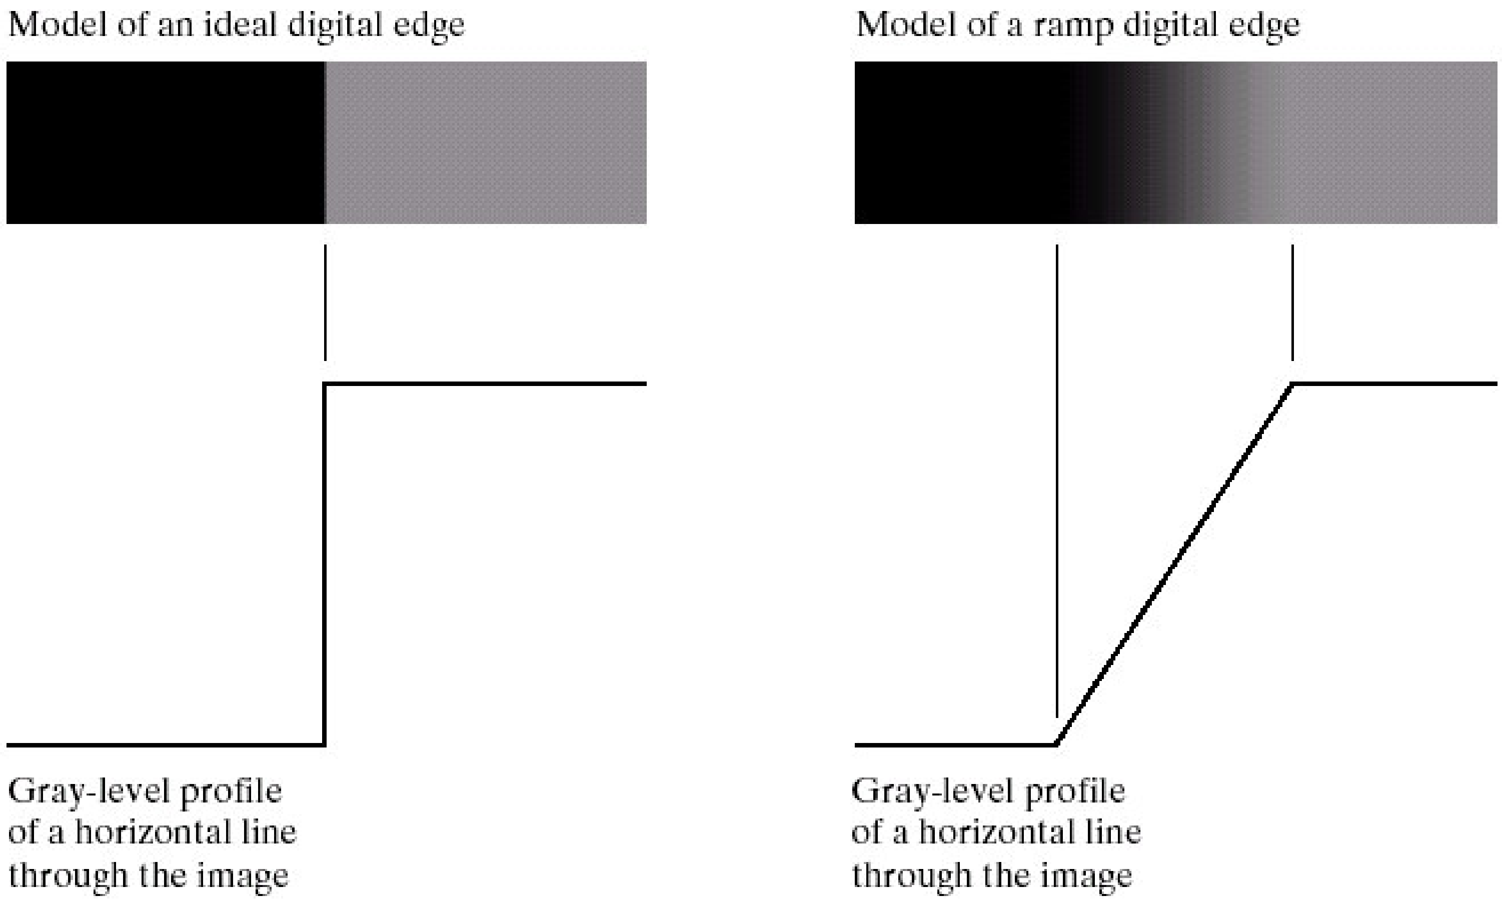

FIGURE 10.5
(a) Model of an ideal digital edge. (b) Model of a ramp edge. The slope of the ramp is proportional to the degree of blurring in the edge.

Line derivatives

a b

FIGURE 10.6
(a) Two regions separated by a vertical edge.
(b) Detail near the edge, showing a gray-level profile, and the first and second derivatives of the profile.

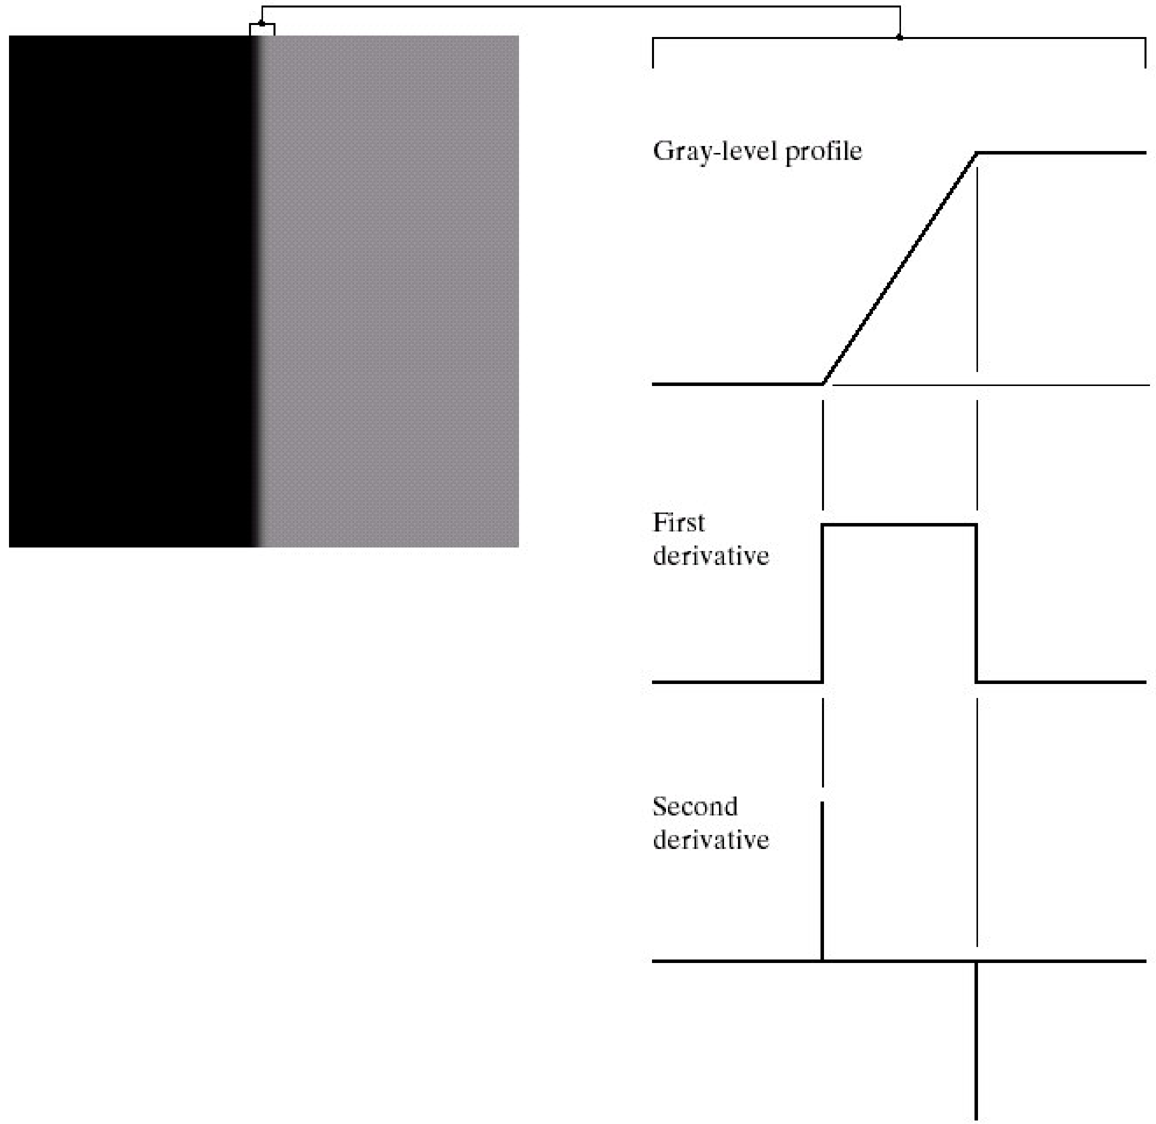

Noisy lines

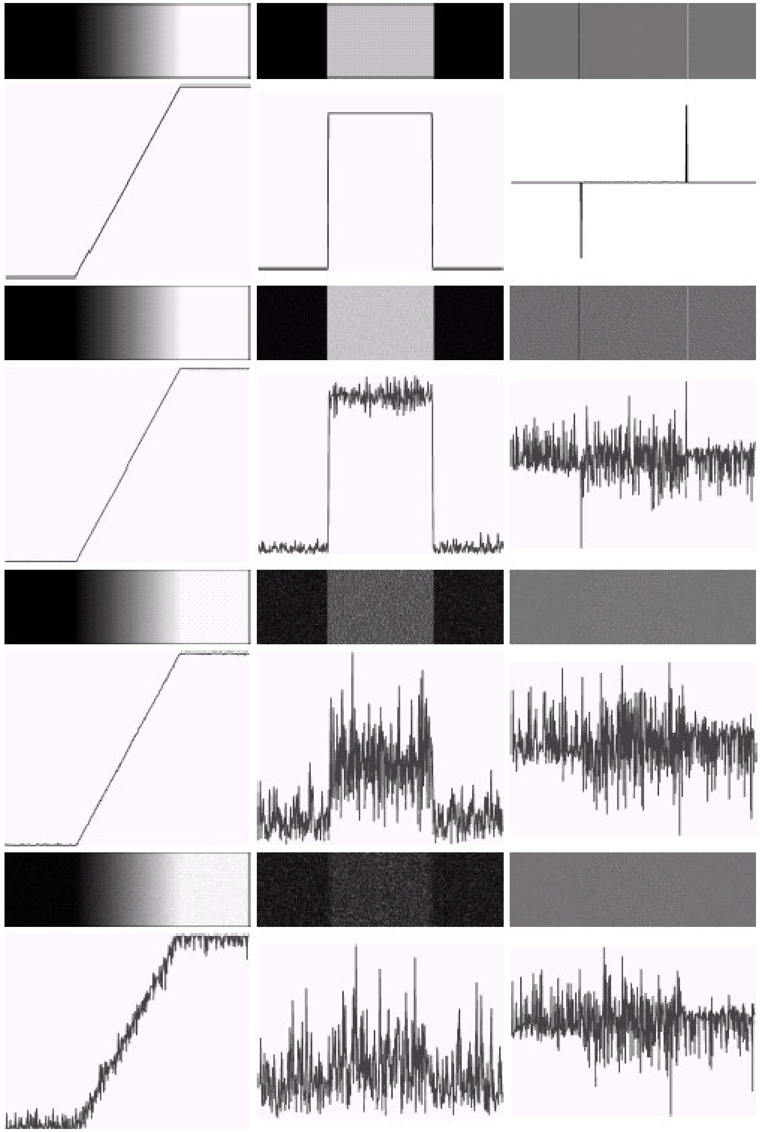

Gradient kernels

a	
b	c
d	e
f	g

FIGURE 10.8 A 3×3 region of an image (the z's are gray-level values) and various masks used to compute the gradient at point labeled z_5.

z_1	z_2	z_3
z_4	z_5	z_6
z_7	z_8	z_9

−1	0
0	1

0	−1
1	0

Roberts

−1	−1	−1
0	0	0
1	1	1

−1	0	1
−1	0	1
−1	0	1

Prewitt

−1	−2	−1
0	0	0
1	2	1

−1	0	1
−2	0	2
−1	0	1

Sobel

Gradient kernels

0	1	1
−1	0	1
−1	−1	0

−1	−1	0
−1	0	1
0	1	1

Prewitt

0	1	2
−1	0	1
−2	−1	0

−2	−1	0
−1	0	1
0	1	2

Sobel

a	b
c	d

FIGURE 10.9 Prewitt and Sobel masks for detecting diagonal edges.

Gradient Example

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, $|G_x| + |G_y|$.

Gradient example after smoothing

a b
c d

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

Laplacian and edge detection

FIGURE 10.13 Laplacian masks used to implement Eqs. (10.1-14) and (10.1-15), respectively.

0	−1	0
−1	4	−1
0	−1	0

−1	−1	−1
−1	8	−1
−1	−1	−1

LoG filter

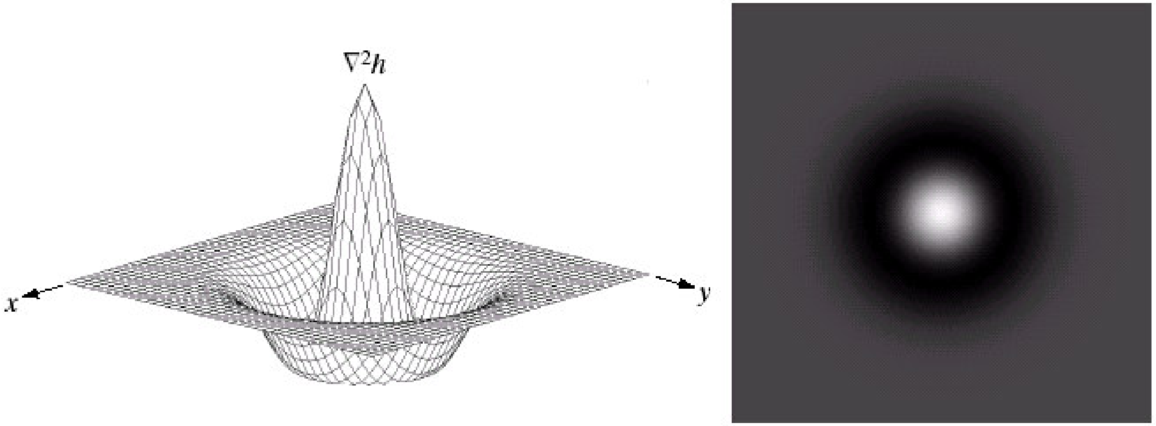

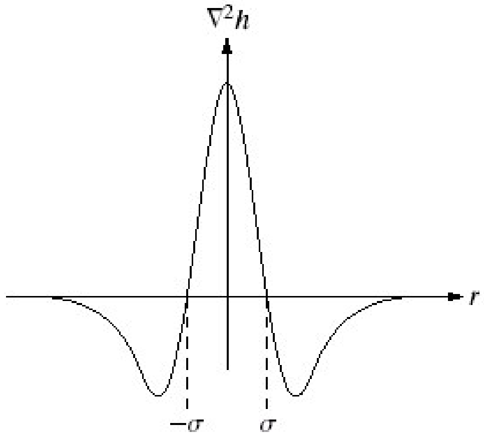

0	0	–1	0	0
0	–1	–2	–1	0
–1	–2	16	–2	–1
0	–1	–2	–1	0
0	0	–1	0	0

a b
c d

FIGURE 10.14 Laplacian of a Gaussian (LoG). (a) 3-D plot. (b) Image (black is negative, gray is the zero plane, and white is positive). (c) Cross section showing zero crossings. (d) 5 × 5 mask approximation to the shape of (a).

LoG example

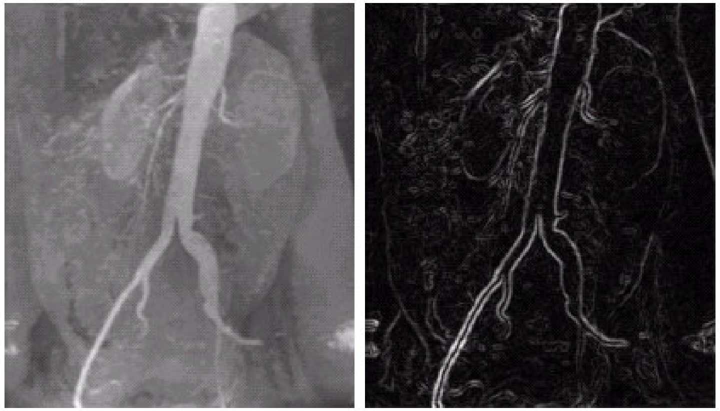

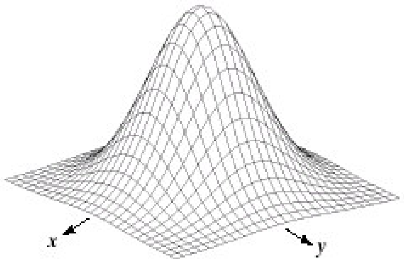

–1	–1	–1
–1	8	–1
–1	–1	–1

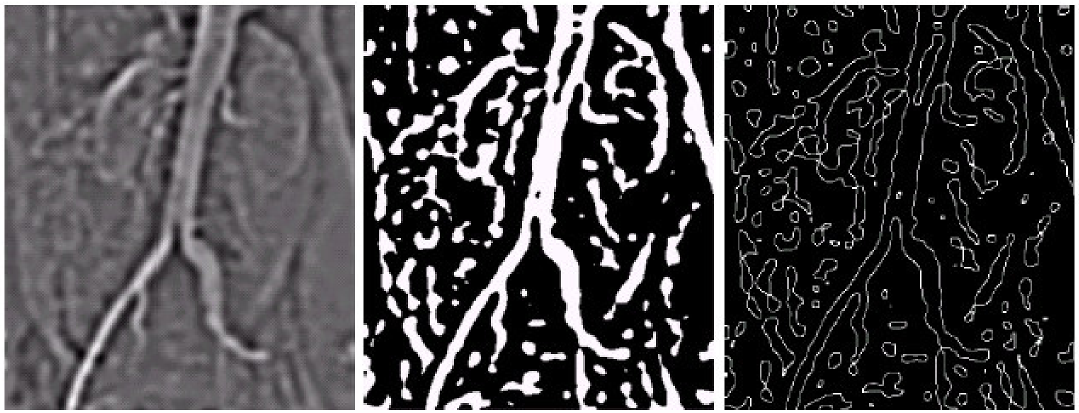

Edge linking

a b
c d

FIGURE 10.16
(a) Input image.
(b) G_y component of the gradient.
(c) G_x component of the gradient.
(d) Result of edge linking. (Courtesy of Perceptics Corporation.)

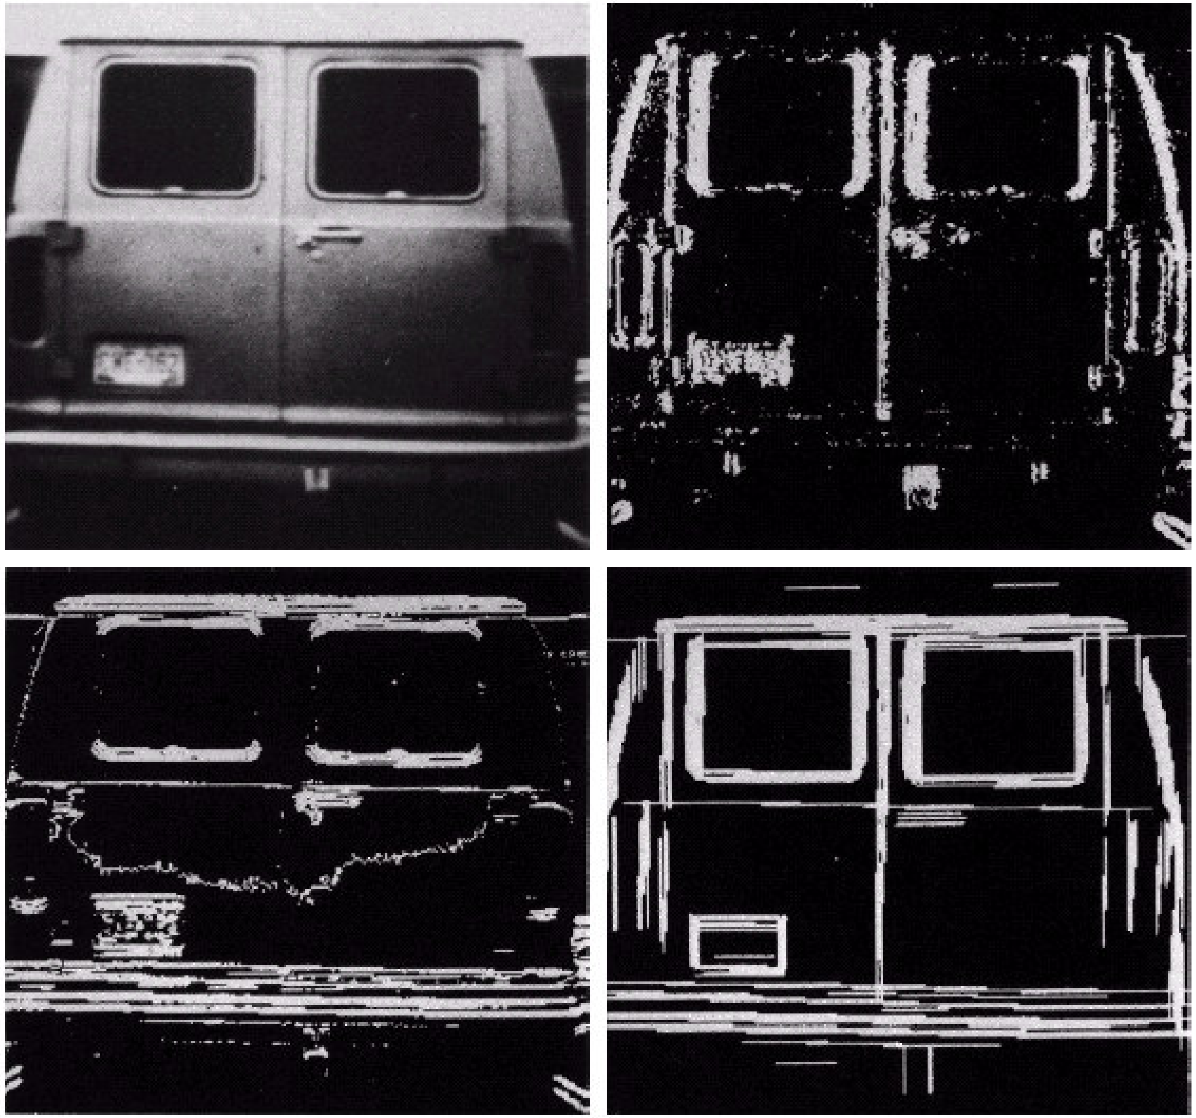

Hough/Radon transform principle

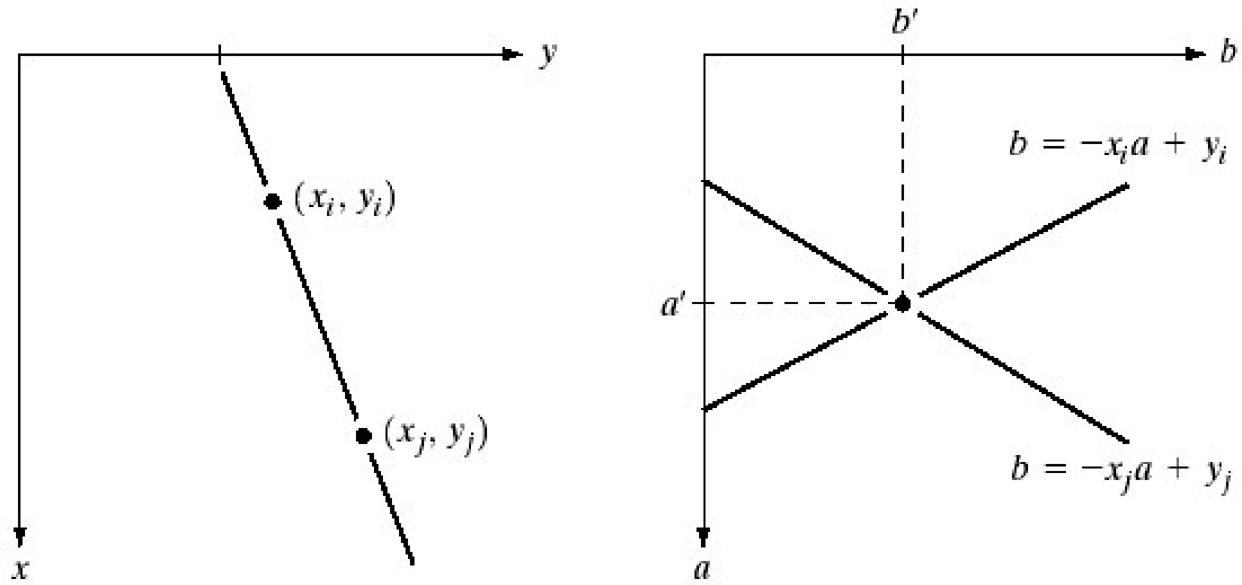

a b

FIGURE 10.17
(a) xy-plane.
(b) Parameter space.

Hough discretization

FIGURE 10.18
Subdivision of the parameter plane for use in the Hough transform.

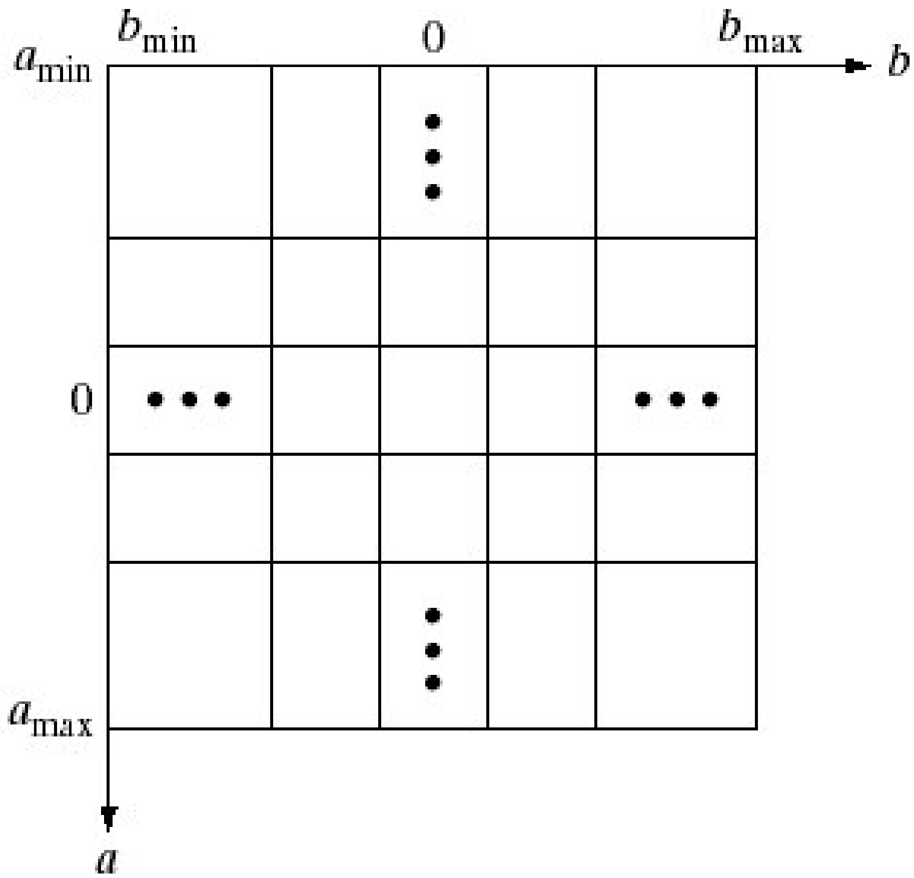

Hough discretization (cont.)

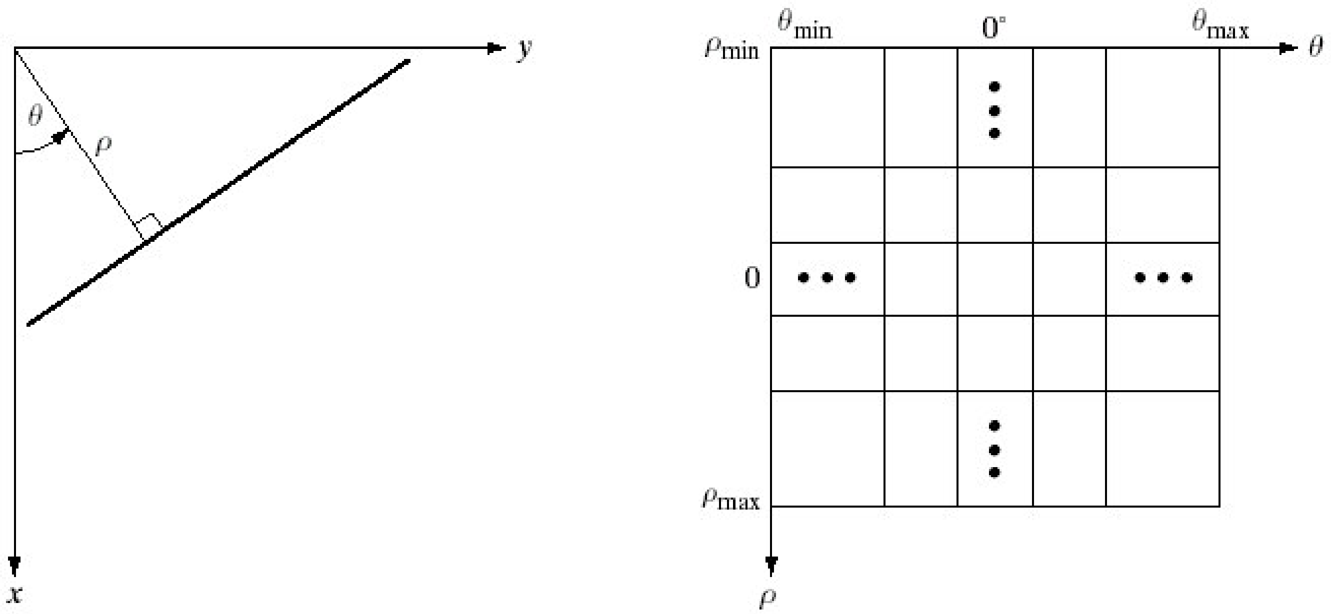

a b

FIGURE 10.19 (a) Normal representation of a line. (b) Subdivision of the $\rho\theta$-plane into cells.

Hough run

a b
c d

FIGURE 10.20 Illustration of the Hough transform. (Courtesy of Mr. D. R. Cate, Texas Instruments, Inc.)

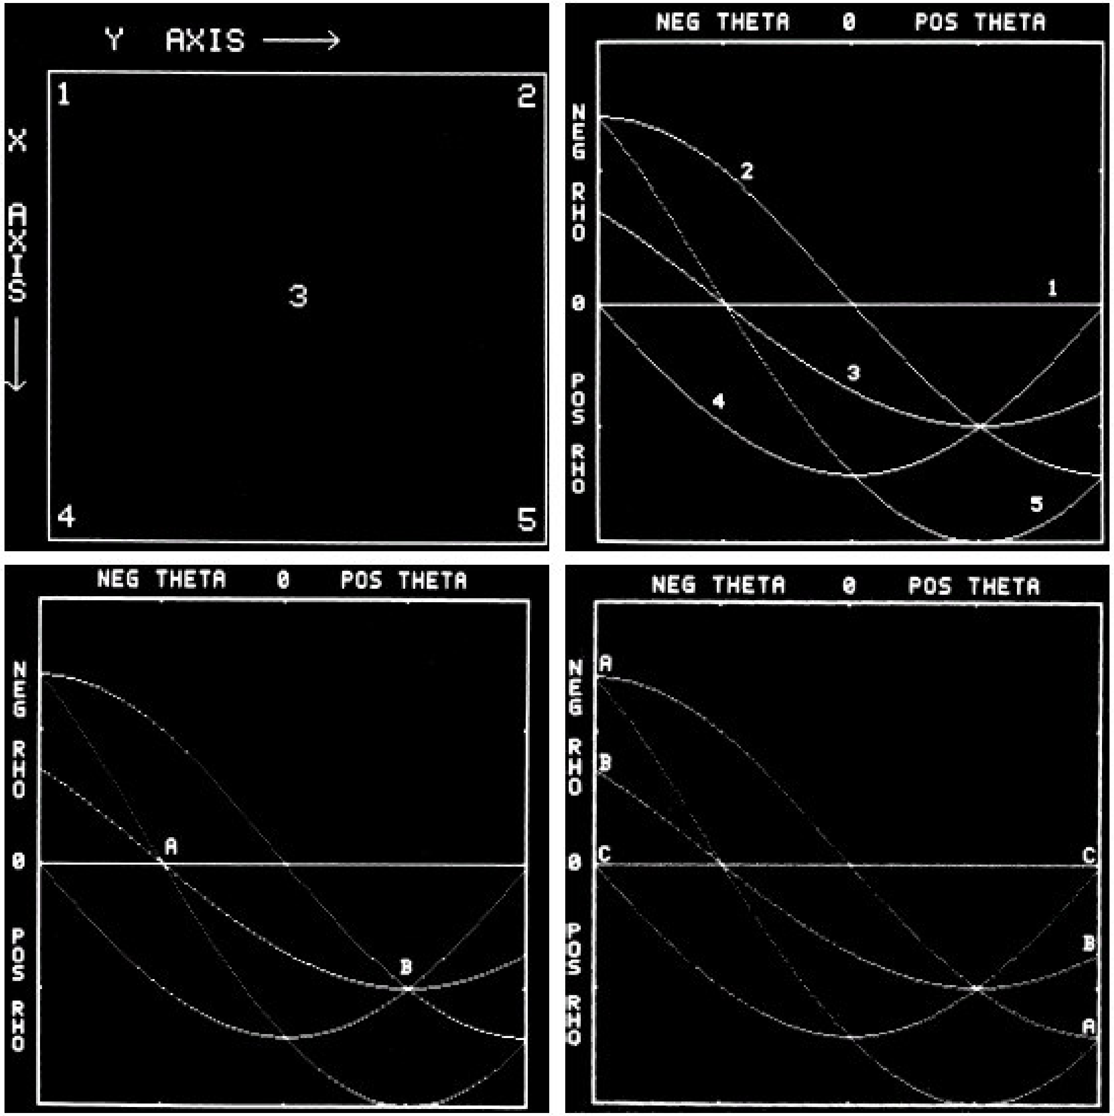

Hough result

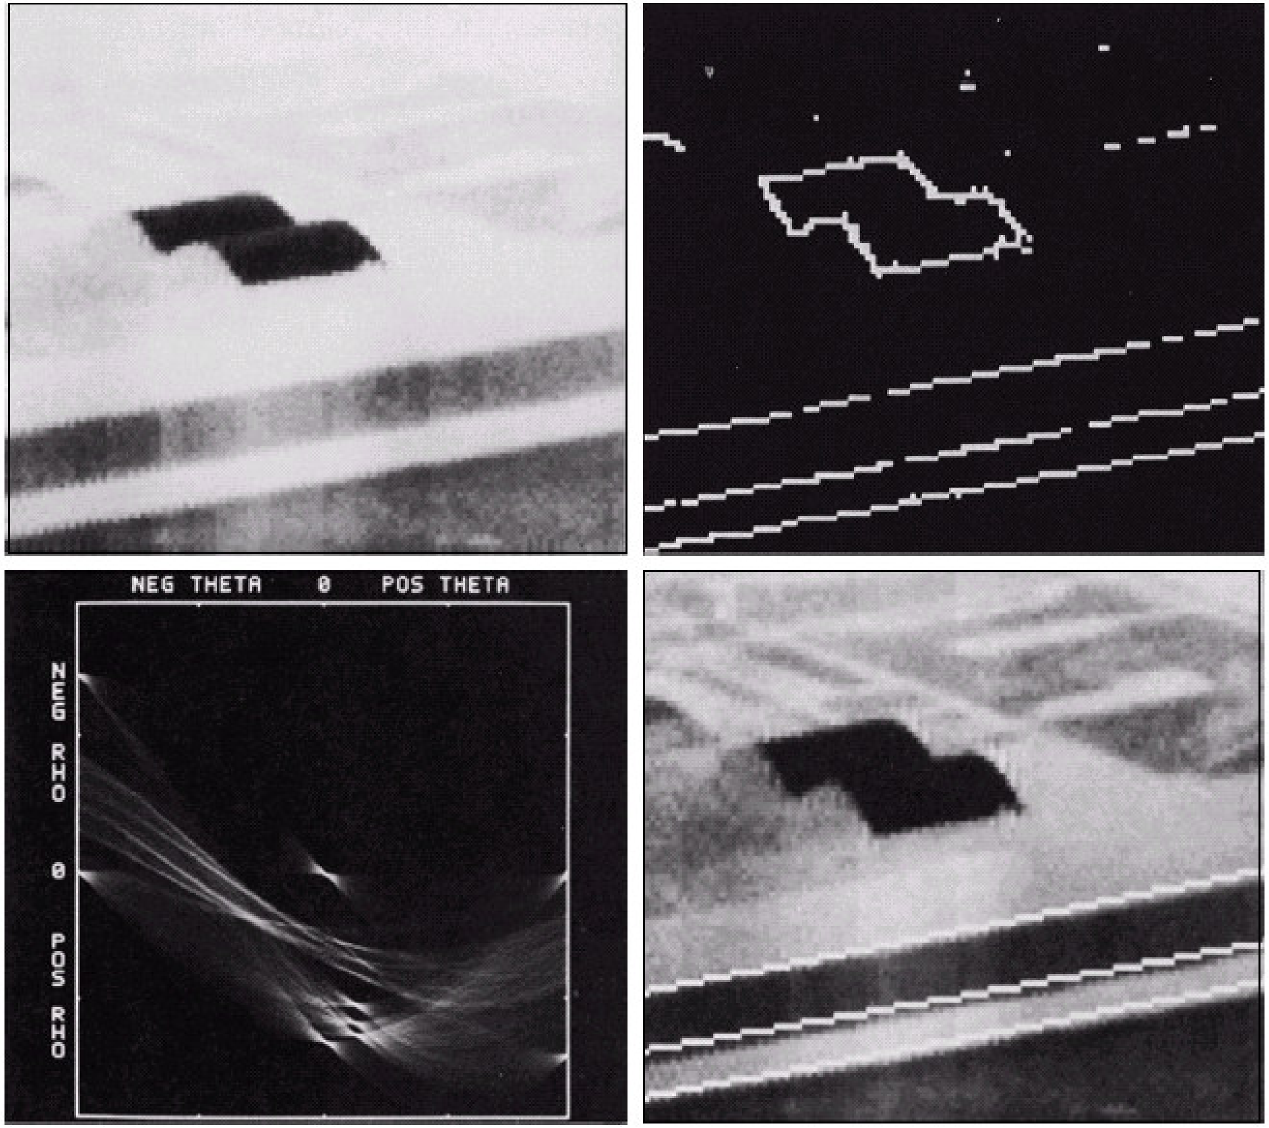

a b
c d

FIGURE 10.21 (a) Infrared image. (b) Thresholded gradient image. (c) Hough transform. (d) Linked pixels. (Courtesy of Mr. D. R. Cate, Texas Instruments, Inc.)

Shortest paths: graph

FIGURE 10.22
Edge element between pixels p and q.

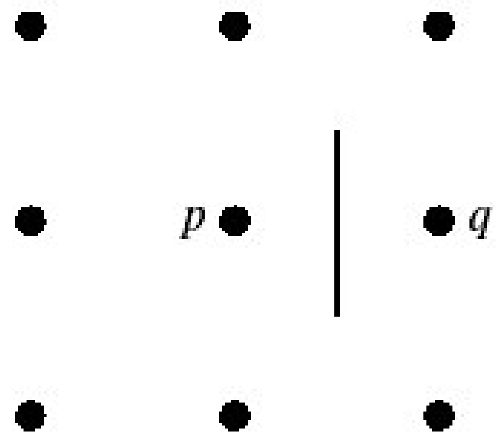

Shortest paths: path

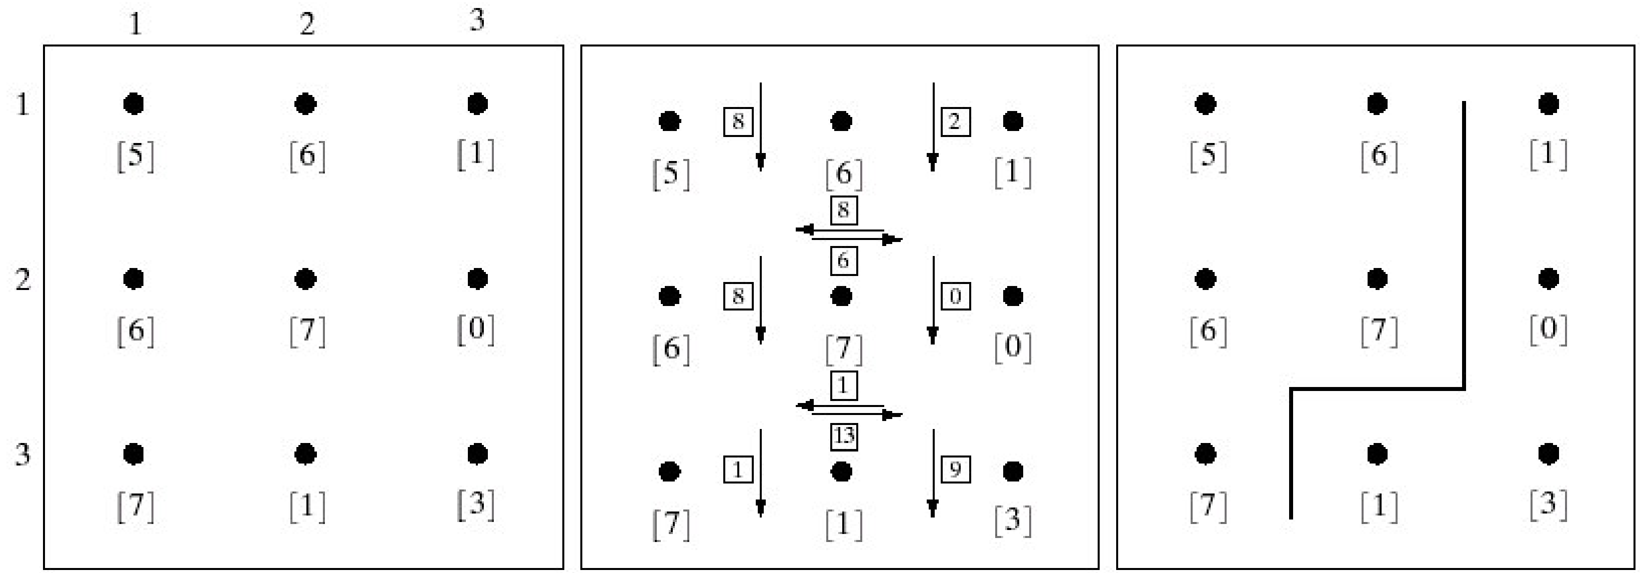

a b c

FIGURE 10.23 (a) A 3 × 3 image region. (b) Edge segments and their costs. (c) Edge corresponding to the lowest-cost path in the graph shown in Fig. 10.24.

Graph search

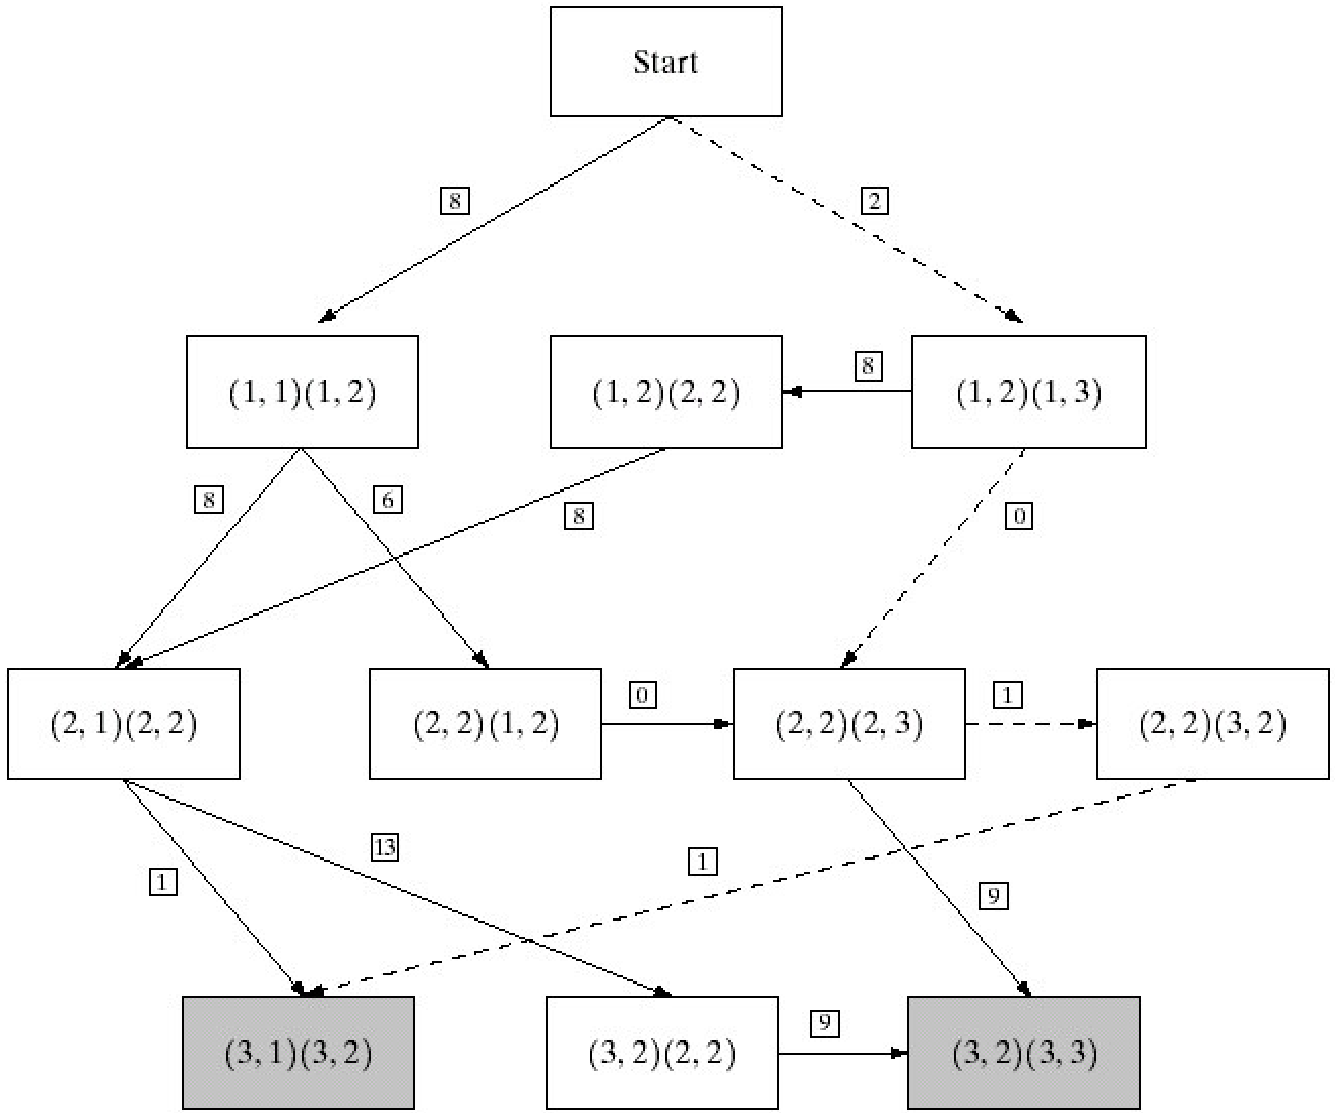

FIGURE 10.24 Graph for the image in Fig. 10.23(a). The lowest-cost path is shown dashed.

Graph search result

FIGURE 10.25 Image of noisy chromosome silhouette and edge boundary (in white) determined by graph search.

Canny edge detection

- multi-stage process
- image is smoothed by Gaussian convolution
- 2-D first derivative operator
- Edges give rise to ridges in the gradient magnitude image
- Algorithm tracks along the top of these ridges
- sets to zero all pixels that are not actually on the ridge top so as to give a thin line in the output = *non-maximum suppression*.
- Tracking proceeds using hysteresis criterion using two thresholds: T_1 and T_2 with $T_1 > T_2$:
  - Tracking can only begin at a point on a ridge higher than T_1
  - Tracking then continues in both directions out from that point until the height of the ridge falls below T_2

Canny examples 1


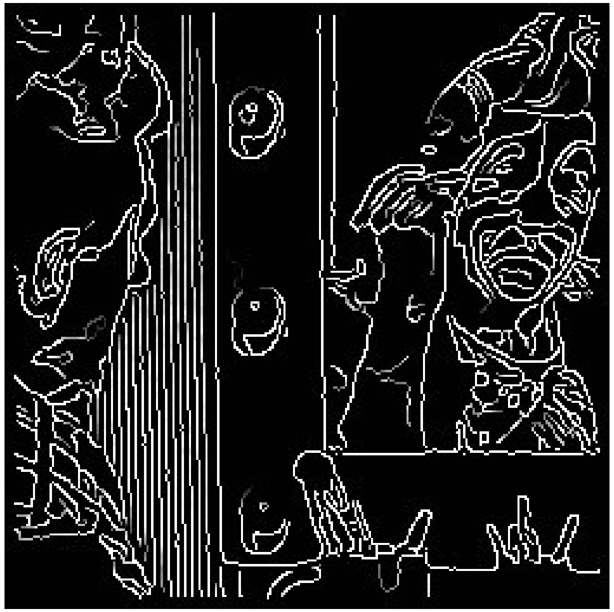


Canny examples 2

 

Summary of edge detection

- Classic technique
- Suitable for many tasks
- Difficult to handle in complex scences
- Not good at delineating regions

IV: region-based segmentation

Morphological segmentation

- Hard, essential, unsolved problem (possibly unsolvable without strong AI), central to image analysis.
- Two main approaches to segmentation: top-down and bottom-up.
- Most literature on segmentation is top-down as a matter of course.
- Mathematical morphology approach is mostly bottom-up.

The watershed line 1

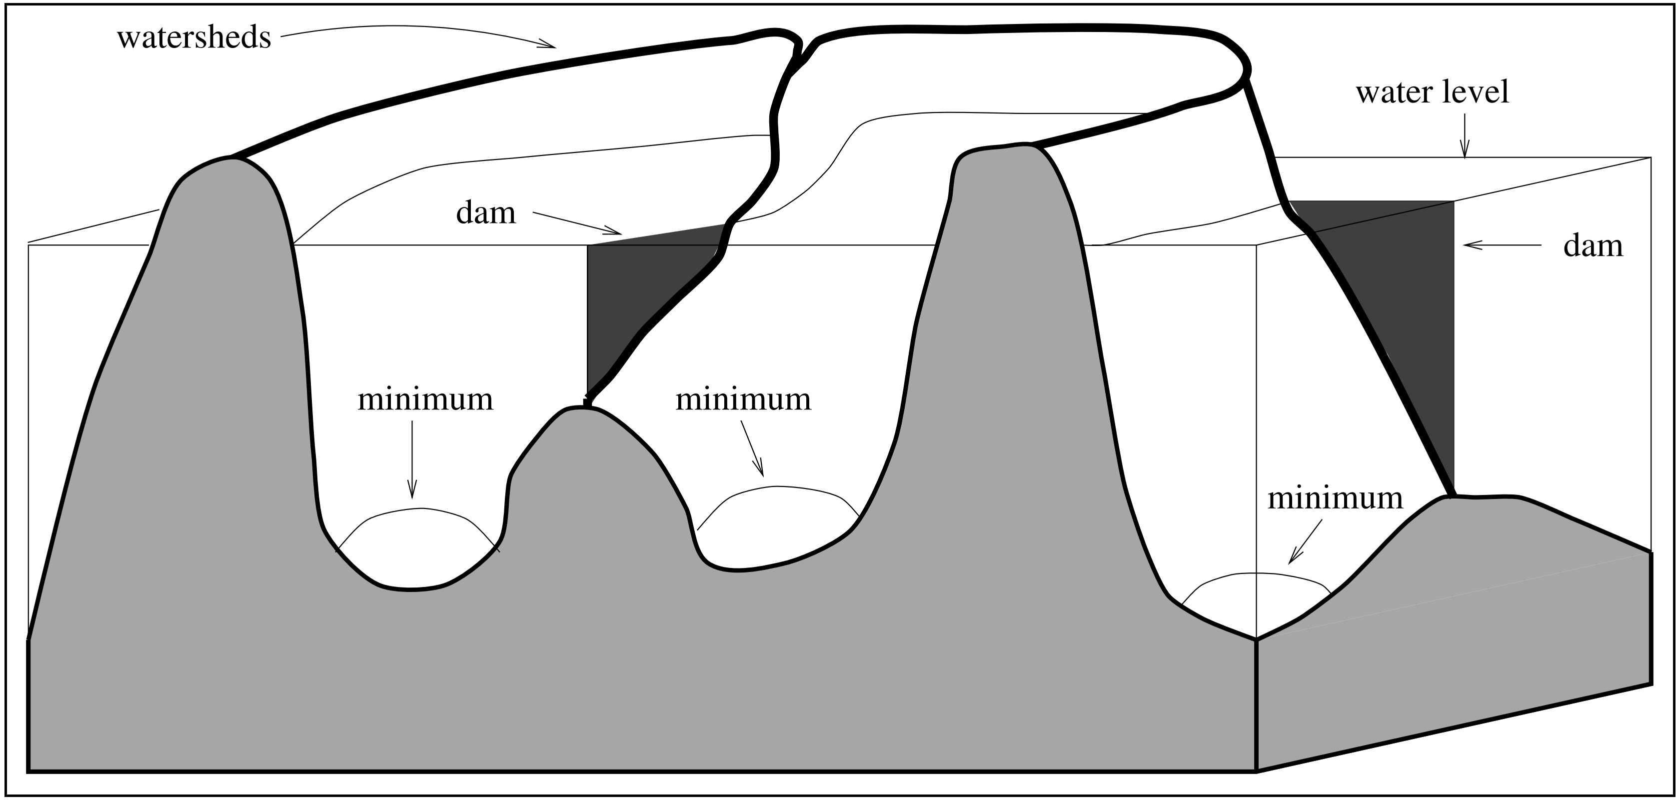

The watershed line 2

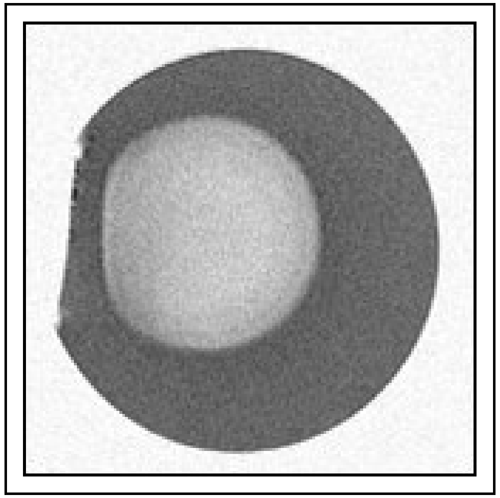

Original

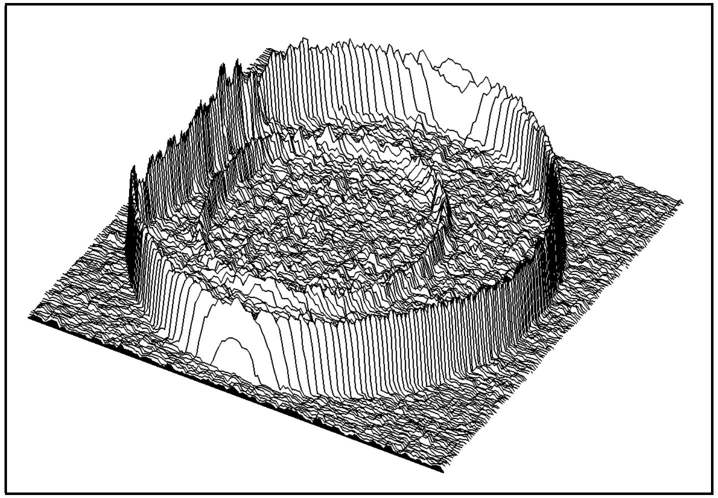

Gradient

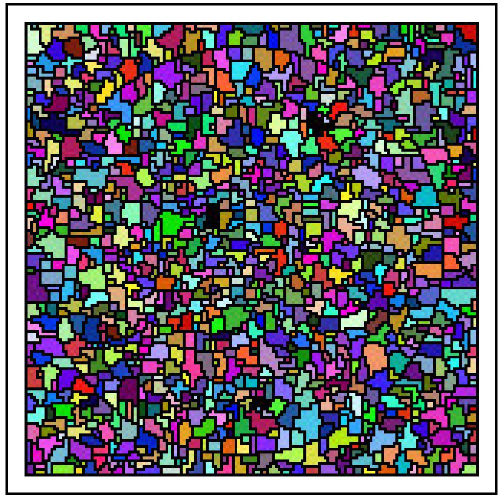

First result

The watershed line 3

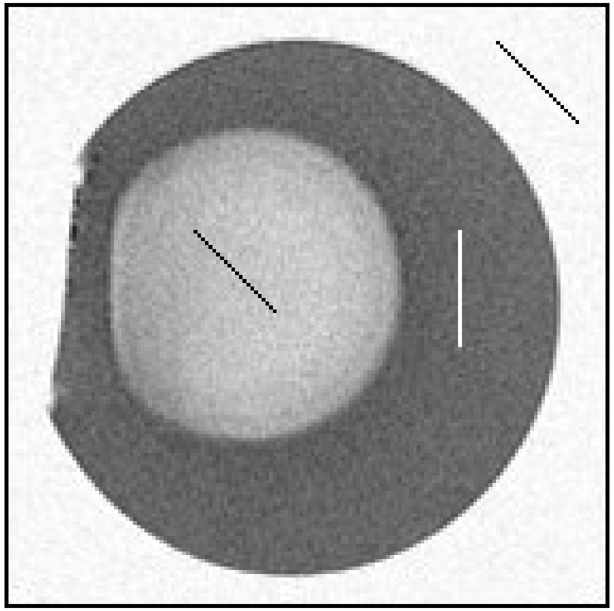

Orig. + Markers

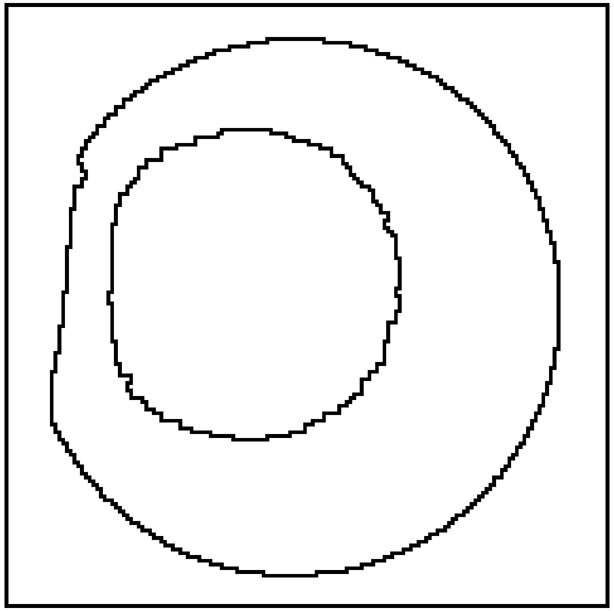

Final

How to use the watershed line?

The classical CMM approach:

- Find a good gradient
- Find good external markers
- Find good internal markers

Other approaches:

- Merge catchment basins afterwards (graph approach f.e.)
- Constrain the watershed

Example of the classical approach (1)

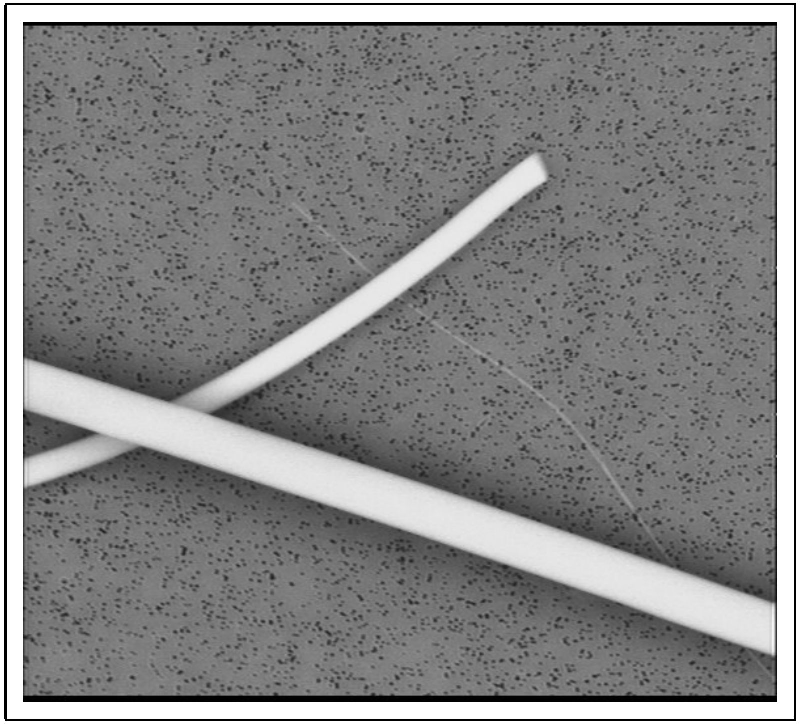

Examples of the classical approach (2)

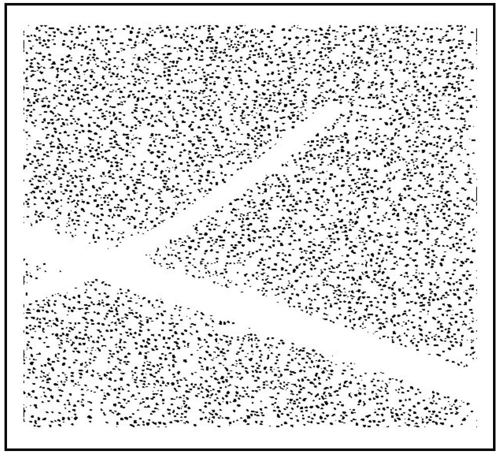

internal

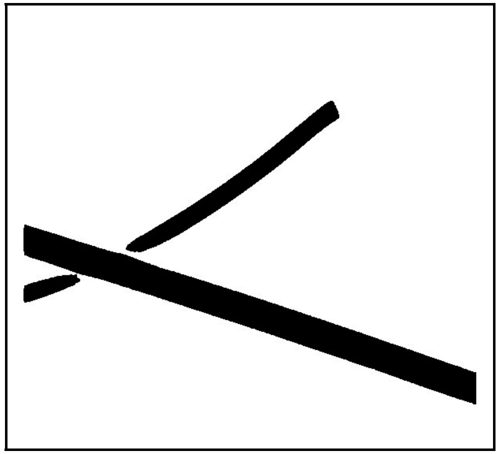

external

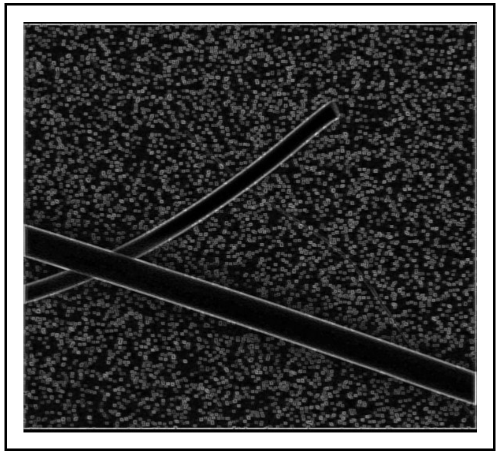

gradient

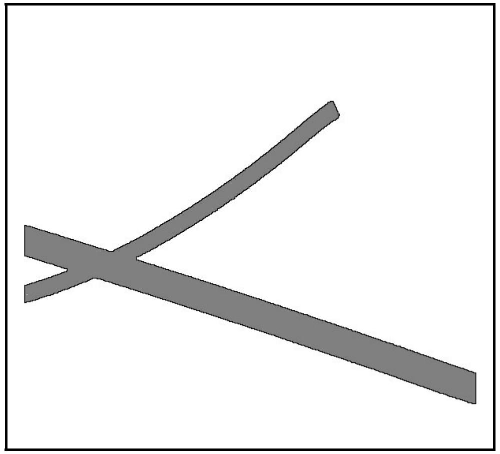

final

Extensions

- Constrained watershed
- Graph-based approaches to reduce oversegmentation
- Applications to image coding

Seeded region growing

How does it work ?

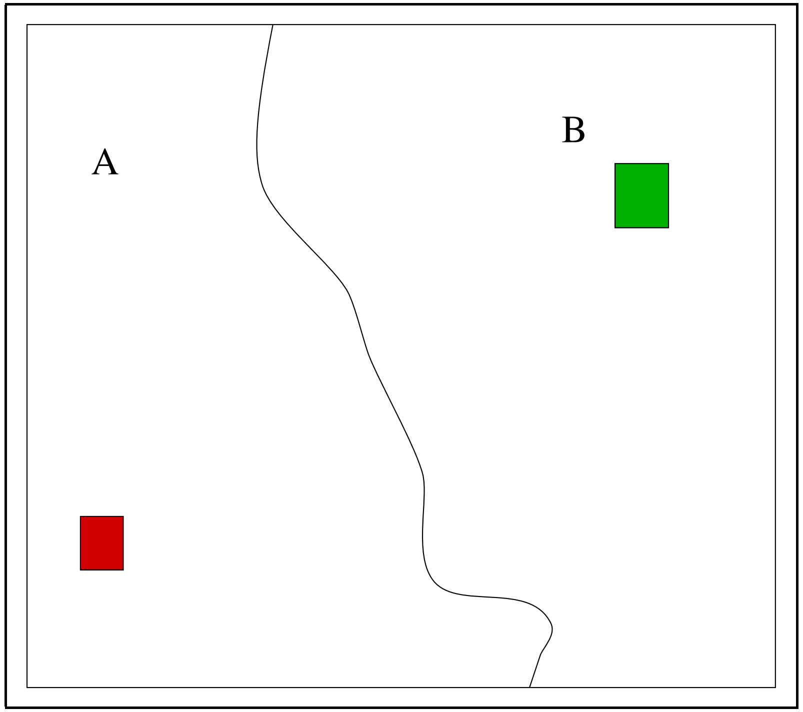

A simple SRG example

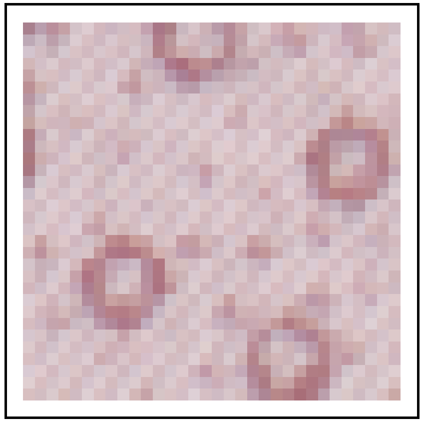
Simple colour image

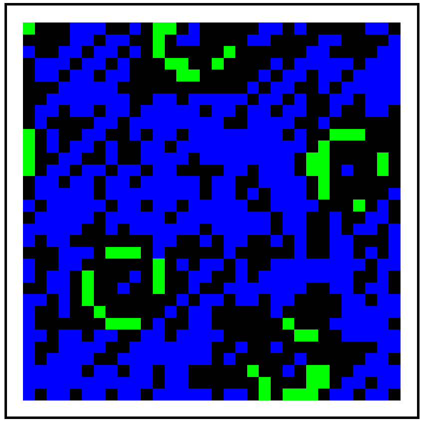
Markers

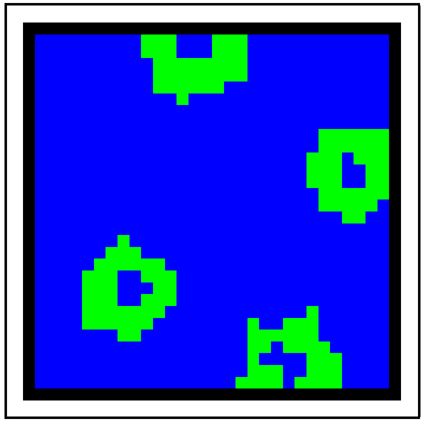
Result

A more complex example

Original

Final

Implementation issues

- In software: priority queues - FIFO w/ priority. Requires random access to the image.
- In hardware: parallel or sequential algorithms. now works in Real Time.

Comparision with other techniques (1)

Fundamental aspects of RBS:

- SRG+WL yield closed contours, like snakes and LS methods.
- Extends natually to 3D.
- Very efficient both in hardware (RT) and software
- SRG+WL are the "mechanical" part of RBS. Intelligent part = constraints, finding internal+external markers
- Usually, some pre-knowledge of the content of the image is necessary (shape, size). Sometimes, markers can be derived from training sets.
- "Ad-hoc" pre-processing.

Comparision with other techniques (2)

Comparison with:

- Thresholding, histogram clustering
- Contour-based methods
- Spatial clustering, Split and Merge
- Active contours
- Level sets
- Knowledge-based methods

An application: microarray analysis

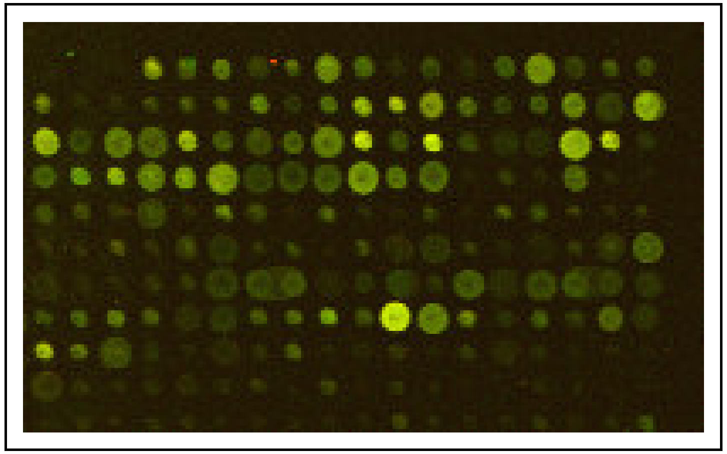

Original

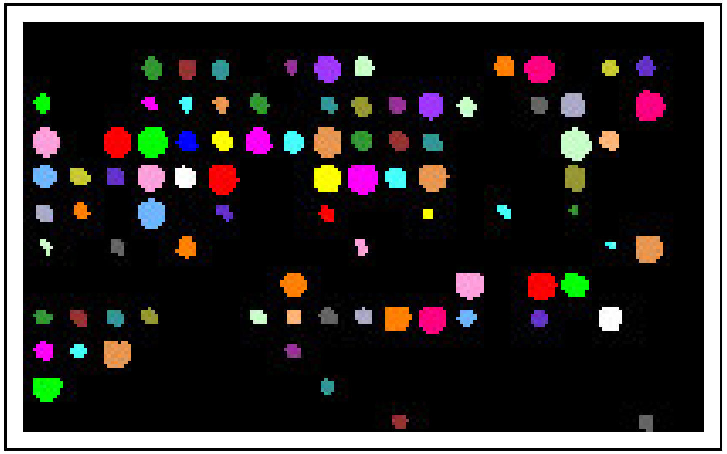

Segmented

Not talked about

- 3D
- Regularization
- Constraints

Lung segmentation

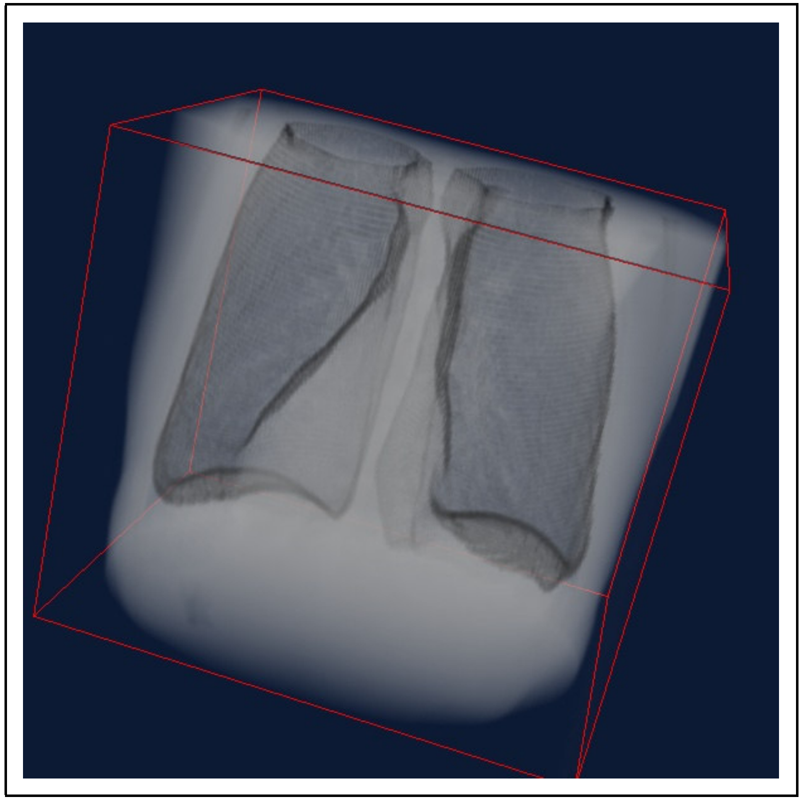

Lung segmentation

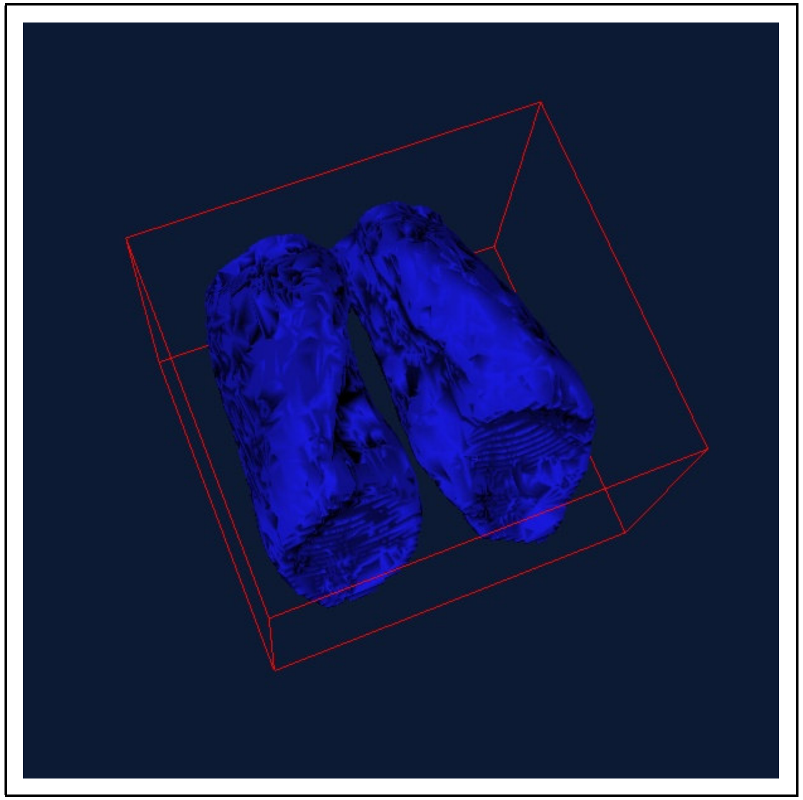

Regularization problems

slice 7

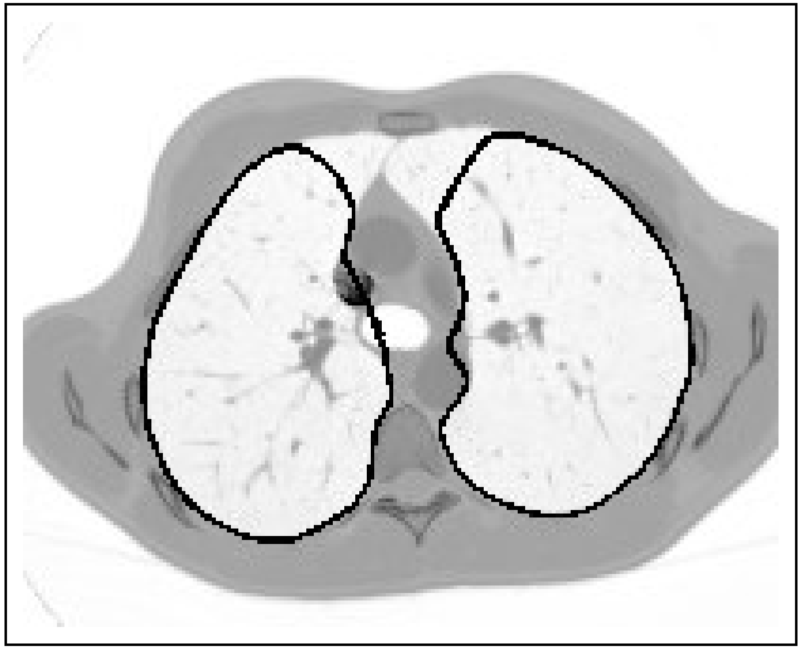

slice 20

Bibliography

[1] G. Matheron. Random sets and integral geometry, 1975, Wiley, New York.

[2] J. Serra. Image Analysis and Mathematical Morphology. 1981, Ass. Press, London.

[3] P. Soille. Morphological Image Analysis – principles and applications. 1999, Springer, Berlin.