Image segmentation

Context

- Segmentation decomposes the image into parts for further analysis
  - Example: background subtraction in human motion analysis
- Once the region of interest is segmented, the representation space can be changed (from image-space to feature space)

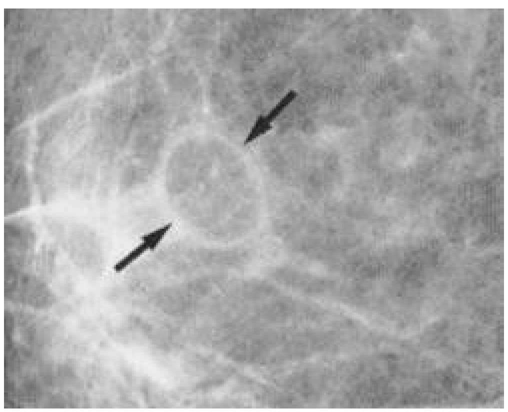

Circumscribed (benign) lesions in digital mammography

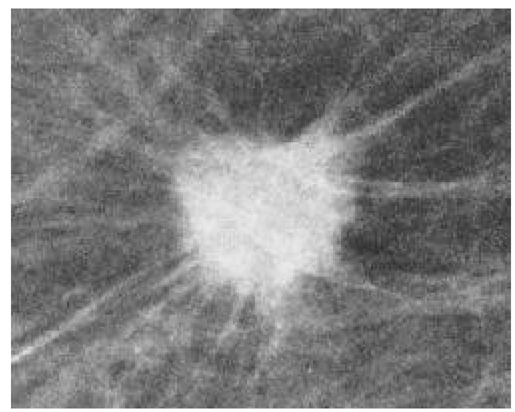

Spiculated (malignant) lesions in digital mammography

What is segmentation ?

- Partitioning an image into regions corresponding to objects
- All pixels in a region share a common property
- Simplest property that pixels can share: intensity
- Thresholding=separation of light and dark regions

A classification of segmentation techniques

- Intensity-based segmentation: Thresholding
- Edge-based segmentation
- Region-based segmentation

A classification of segmentation techniques

- Intensity-based segmentation: Thresholding
- Edge-based segmentation
- Region-based segmentation

Assumptions for thresholding

1. the intensity values are different in different regions
2. within each region, which represents the corresponding object in a scene, the intensity values are similar.

Original

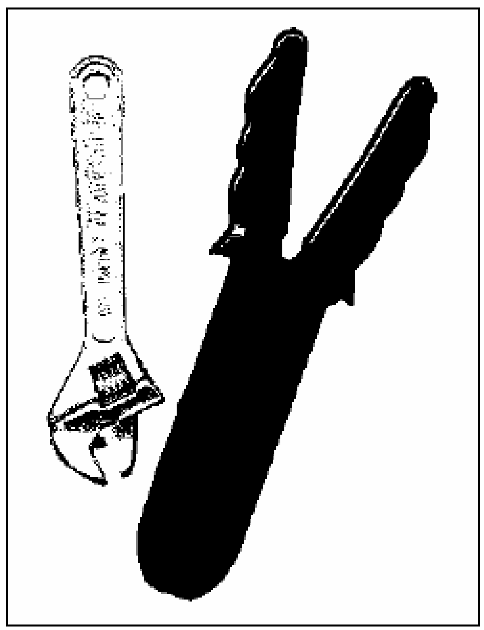

Threshold = 50

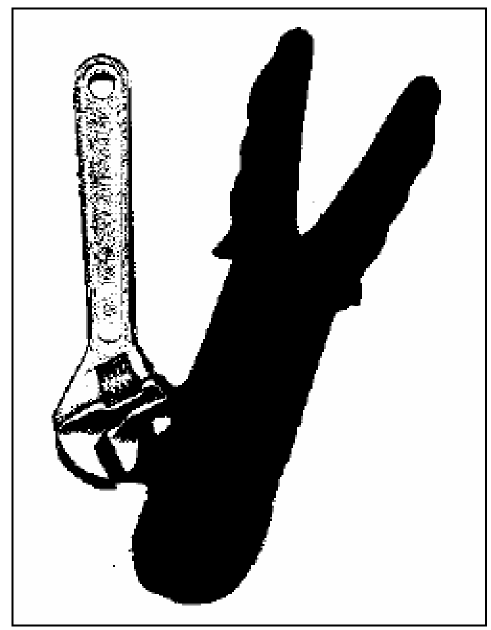

Threshold = 75

Intensity-based thresholding

Image thresholding classifies pixels into two categories:
– Those to which some property measured from the image falls below a threshold, and those at which the property equals or exceeds a threshold.
– Thresholding creates a binary image : binarization e.g. perform cell counts in histological images

Choosing a threshold is a critical task.

```
n=imread('nodules1.jpg');
figure(1); imshow(n);
n1=im2bw(n,0.35);
n2=im2bw(n,0.75);
figure(2), imshow(n1);
figure(3), imshow(n2);
```

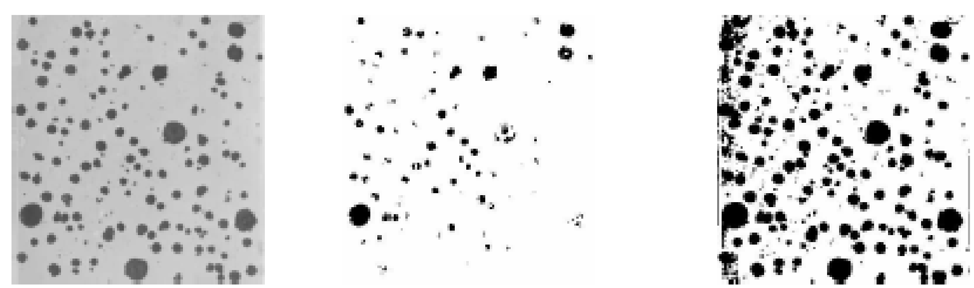

n: Original image n1: Threshold too low n2: Threshold too high

Fixed versus dynamic thresholding

- In fixed (or global) thresholding, the threshold value is held constant throughout the image:

$$g(x,y) = \begin{cases} 0 & f(x,y) < T \\ 1 & f(x,y) >= T \end{cases}$$

- Local (or dynamic thresholding): depends on the position in the image. The image is divided into overlapping sections which are thresholded one by one.

Threshold detection methods

- P-tile thresholding
- Optimal thresholding
- Mixture modelling
- Adaptive thresholding

- Note: all the above methods are automated

P-tile method

- a priori information: object is brighter/darker than background and occupies a certain *known* percentile 1/p from the total image area (example: printed text sheet)
- We set the threshold by finding the intensity level such that 1/p image pixels are below this value
- We use the cumulative histogram

$$c(g) = \sum_{k=0}^{g} h(k)$$

$$h(k) = \frac{n_k}{n}$$

- *T* verifies the equation $c(T)$=1/p (for a dark foreground)
- $c(T)$=1-1/p (for a bright foreground)

Finding modes

- Histogram shape analysis
- Foreground pixels form one peak
- Background pixels form the second peak
- Intuitively: the threshold is set as the gray level that has a minimum value between two maxima
- Problem: noisy histograms (salt-and pepper noise)

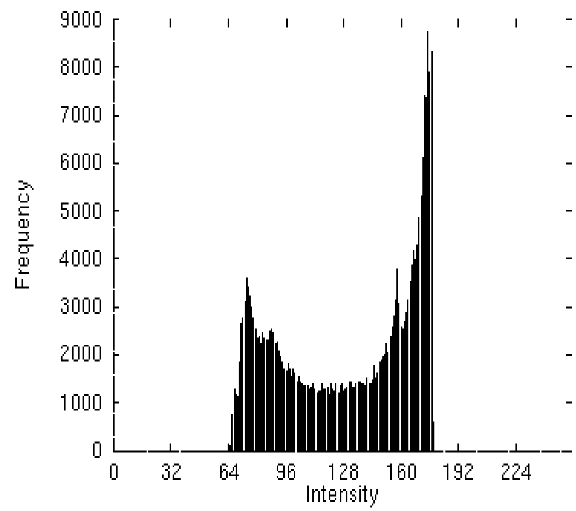

Figure 5.3 *A bimodal histogram.*

Optimal thresholding

- Idea: the histogram of an image is the sum of two overlapping distributions
- Optimal threshold: overlapping point of these distributions (corresponds to the minimum probability between the maxima of 2 distributions)
- Problem: distributions are unknown

Comparison between conventional and optimal thresholding

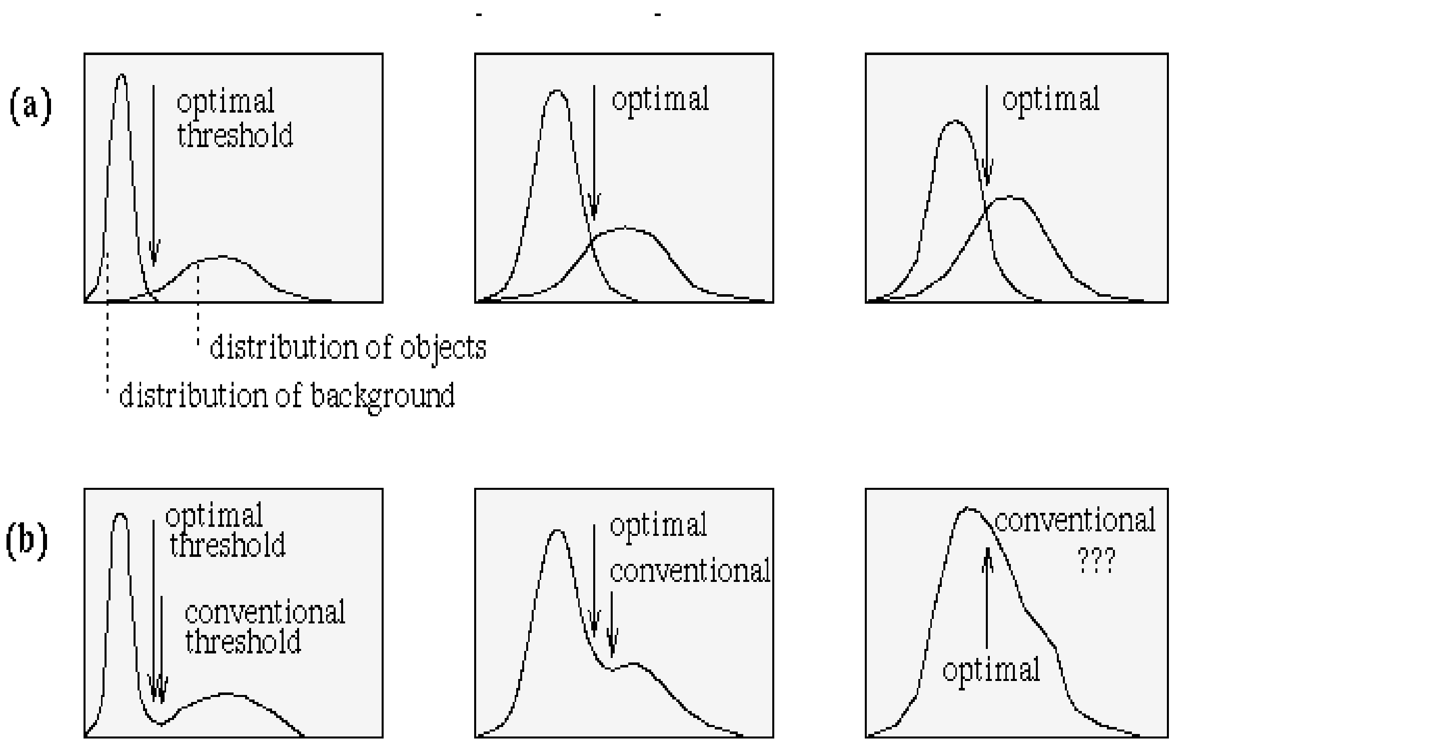

Figure 5.4 *Grey level histograms approximated by two normal distributions; the threshold is set to give minimum probability of segmentation error: (a) Probability distributions of background and objects, (b) corresponding histograms and optimal threshold.*

Optimal thresholding by clustering

- Simplest case: segmentation into two classes (object/background).
- The intensities in each class will be our clusters.
- We want to find a threshold such that each pixel on each side of the threshold is closer in intensity to the mean of all pixels on that side of the threshold than to the mean of all pixels on the other side of the threshold.

Iterative optimal threshold selection

1. Select an initial estimate for T

2. Segment the image using T. This produces 2 groups: G1 pixels with value >T and G2 , with value <T

3. Compute $\mu 1$ and $\mu 2$, average pixel values of G1 and G2

4. New threshold: $T=1/2(\mu 1+\mu 2)$

5. Repeat steps 2 to 4 until T stabilizes.

Iterative Clustering Algorithm

m1(1)=260.83, m2(1)=539.00
m1(2)=39.37, m2(2)=1045.65
m1(3)=52.29, m2(3)=1098.63
m1(4)=54.71, m2(4)=1106.28
m1(5)=55.04, m2(5)=1107.24
m1(6)=55.10, m2(6)=1107.44
m1(7)=55.10, m2(7)=1107.44

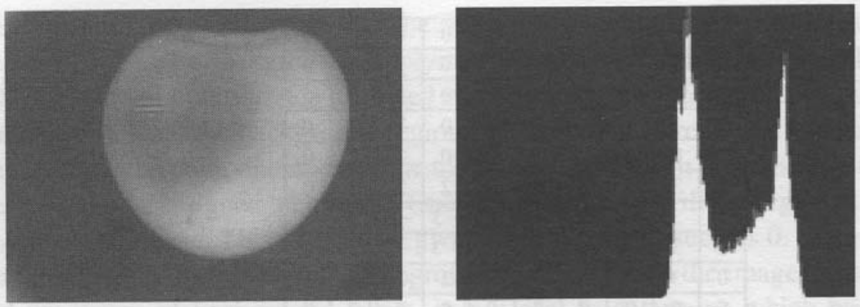

(*a*) Image of a bruised cherry (*b*) Histogram of the cherry image

Figure 3.22 Histogram of the image of a bruised cherry displaying two modes: (*a*) one representing the bruised portion and (*b*) the nonbruised portion. (Courtesy of Patchrawat Uthaisombut.)

Optimal thresholding : the Otsu method

- Optimal thresholding methods select the threshold based on the minimization of a criterion function.
- The criterion for Otsu is the minimization of the within-group variance of the two groups of pixels separated by the threshold.

(*a*) Original image (*b*) Pixels below 93 (*c*) Pixels above 93

Figure 3.24 A gray-tone image and the pixels below and above the threshold of 93 (shown in white) found by the Otsu automatic thresholding operator. (Original image courtesy of John Illingworth and Ata Etamadi.)

Mixture modelling

- Assumption: region intensities are each normal distributions (Gaussians)

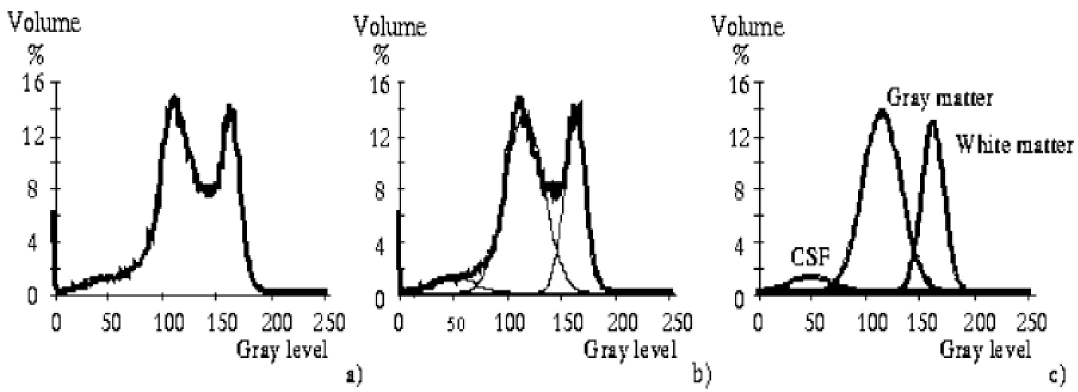

Figure 5.5 *Segmentation of 3-D T1-weighted MR brain image data using optimal thresholding: (a) Local gray level histogram, (b) fitted Gaussian distributions; global 3-D image fit, (c) Gaussian distributions corresponding to WM, GM, and CSF.*

Mixture modelling (cont'd)

- Each of the Gaussian distributions has a mean and standard deviation *independent* of the threshold that we choose
- Foreground/background case:

$$h_{\text{model}}(g) = n_B\ e^{-(g-\mu_B)^2/2\sigma_B^2} + n_O\ e^{-(g-\mu_O)^2/2\sigma_O^2}$$

- We need to estimate 6 parameters
- Evaluation of how well the sum of the distributions approximate the histogram

$$F = \sum_{0}^{N-1} \left[h_{\text{model}}(g) - h_{\text{image}}(g) \right]^2$$

- The parameters will be chosen such as to minimize the error F

Thresholding and illumination

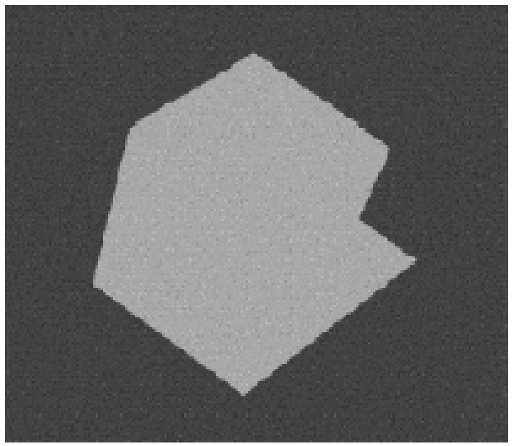

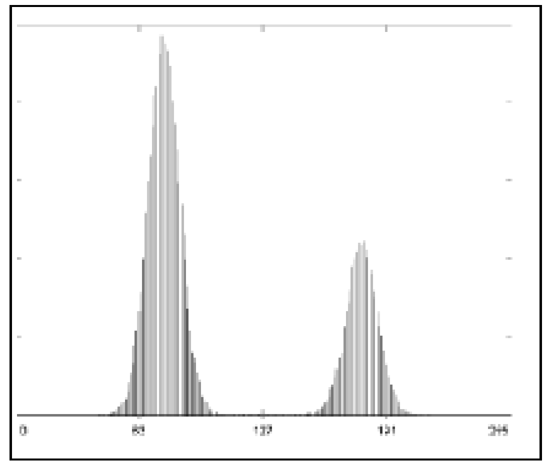

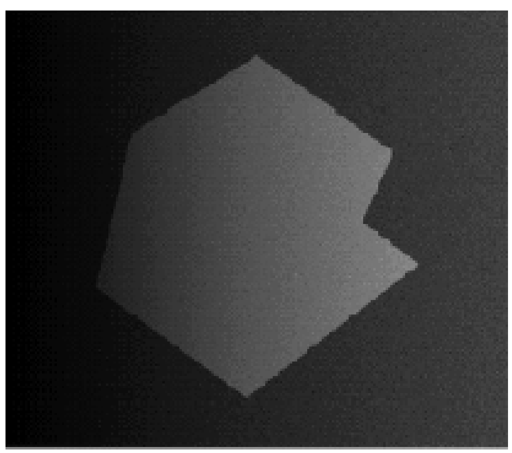

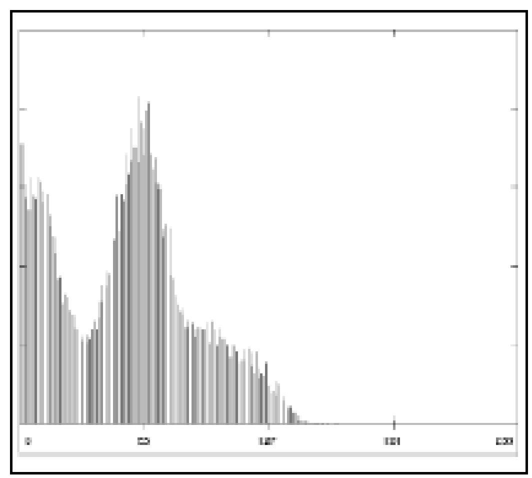

Connected component labeling

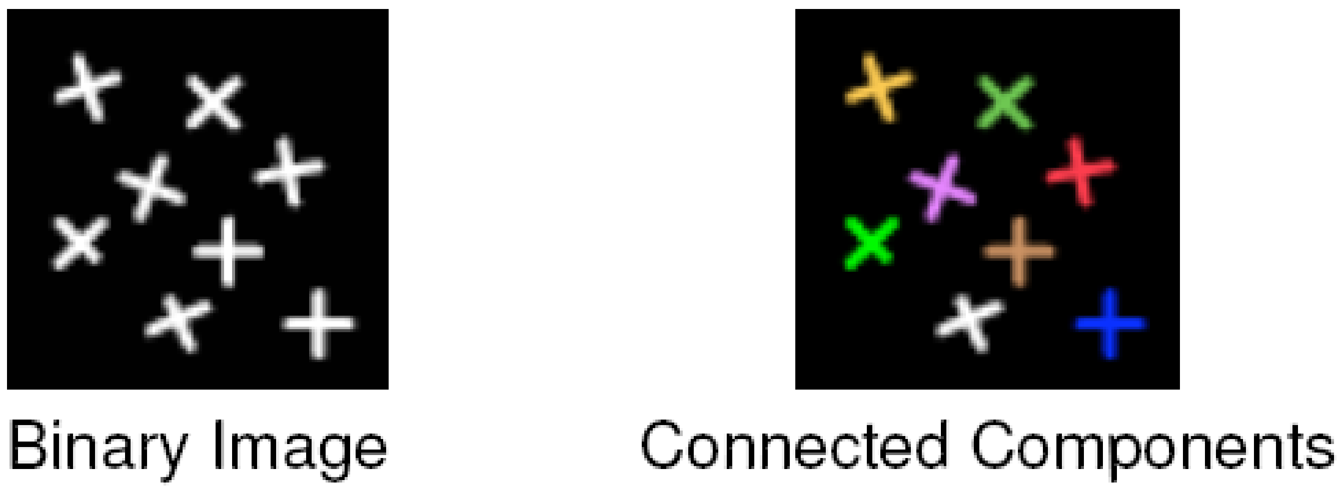

From Brian Morse, http://morse.cs.byu.edu/650/home/index.php

Connected component labeling

- 1. Scan through the image pixel by pixel across each row in order:
  - If the pixel has no connected neighbors with the same value that have already been labeled, create a new unique label and assign it to that pixel.
  - If the pixel has exactly one label among its connected neighbor with the same value that has already been labeled, give it that label.
  - If the pixel has two or more connected neighbors with the same value but different labels, choose one of the labels and remember that these labels are equivalent.
- 2. Resolve the equivalencies by making another pass through the image and labeling each pixel with a unique label for its equivalence class.

From Brian Morse, http://morse.cs.byu.edu/650/home/index.php

Thresholding: Summary

- Advantages:
- Simple to implement
- Fast (especially if repeating on similar images)
- Good for some kinds of images (e.g., documents, controlled lighting)
- Disadvantages:
- No guarantees of object coherency—

may have holes, extraneous pixels, etc.

(incomplete) solution: post-processing with morphological operators

Next lecture

- Edge-based segmentation
- Region-based segmentation