Thresholding

(Jain et al., Sections 3.2.1, 3.2.2, Petrou et al., Chapt 7)

- The simplest approach to segment an image is using thresholding.

$$\text{If } f(x, y) > T \text{ then } f(x, y) = 0 \text{ else } f(x, y) = 255$$

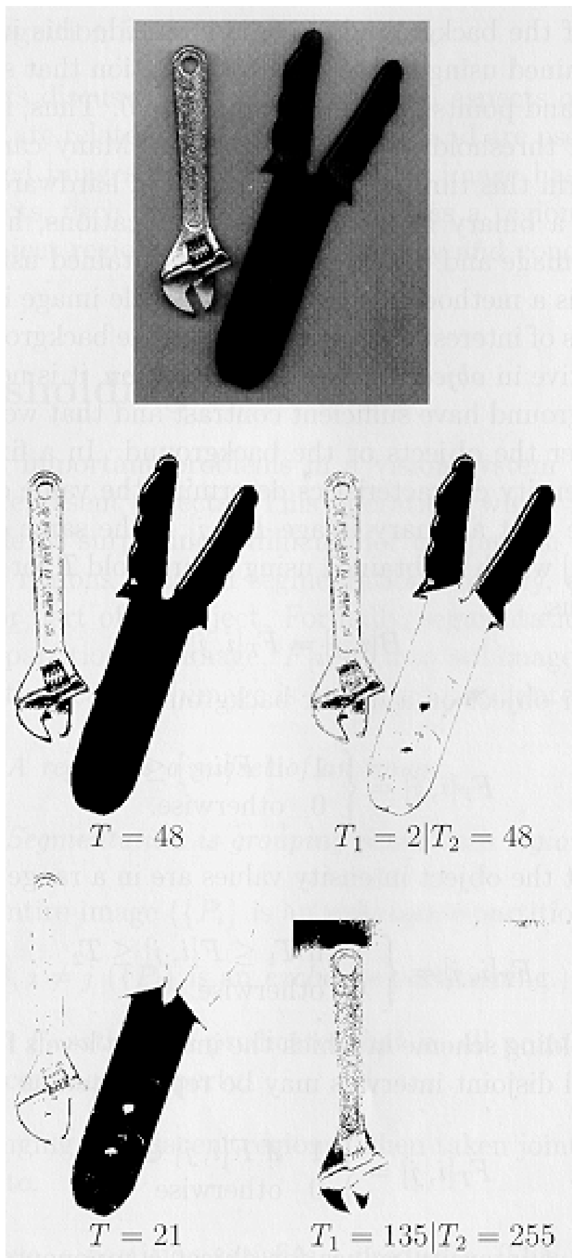

• **Automatic thresholding**

- To make segmentation more robust, the threshold should be automatically selected by the system.

- Knowledge about the objects, the application, the environment should be used to choose the threshold automatically:

  * Intensity characteristics of the objects
  * Sizes of the objects
  * Fractions of an image occupied by the objects
  * Number of different types of objects appearing in an image

• Choosing the theshold using the image histogram

- Regions with uniform intensity give rise to strong peaks in the histogram !!

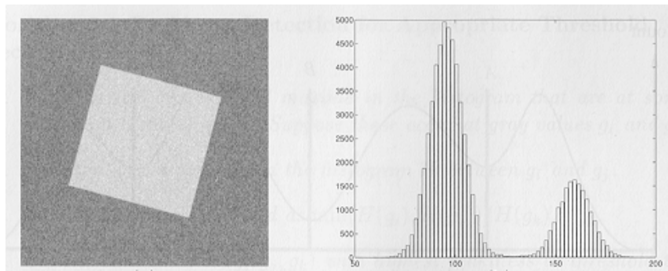

- Multilevel thresholding is also possible (although more difficult in practice)

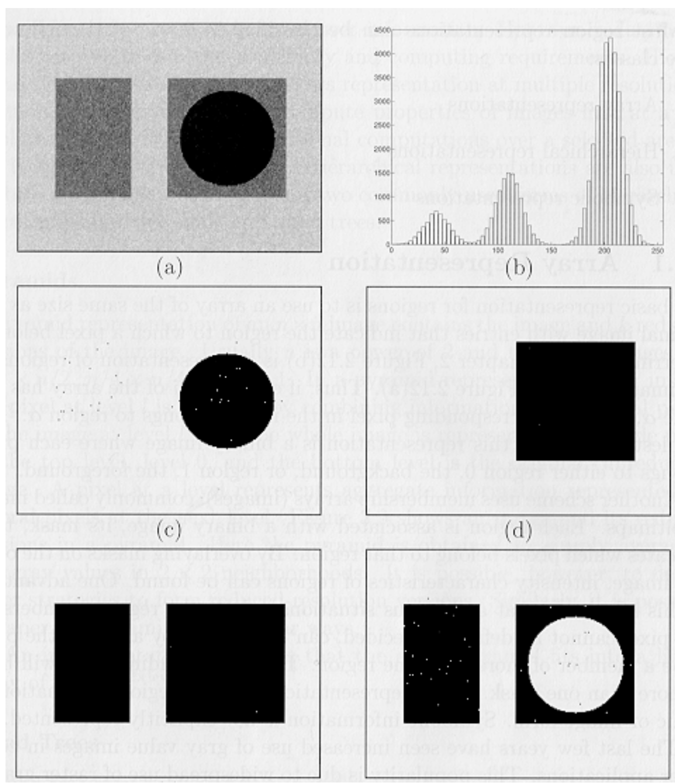

If $f(x, y) < T_1$ then $f(x, y) = 255$

else if $T_1 \leq f(x, y) < T_2$ then $f(x, y) = 128$

else $f(x, y) = 0$

- In general, good thershold can be selected if the histogram peaks are tall, narrow, symmetric, and separated by deep valleys.

• Hysteresis thresholding

- If there is no clear valley in the histogram of an image, it means that there are several background pixels that have similar gray level value with object pixels and vice versa.

- Hystreresis thresholding (i.e., two thresholds, one at each side of the valley) can be used in this case.

- Pixels above the high threshold are classified as object and below the low threshold as background.

- Pixels between the low and high thresholds are classified as object only if they are adjacent to other object pixels.

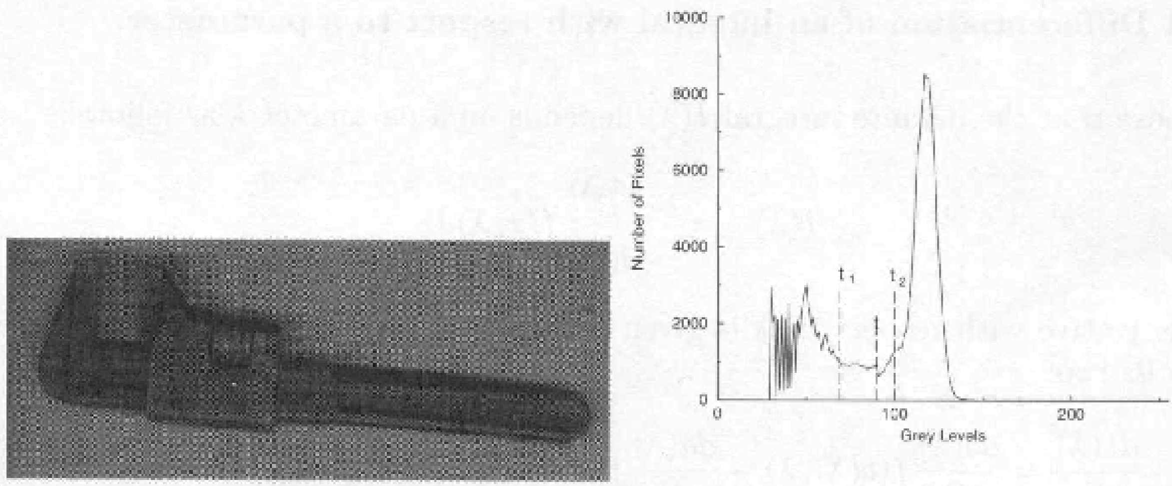

(a) Original image (b) Histogram of (a)

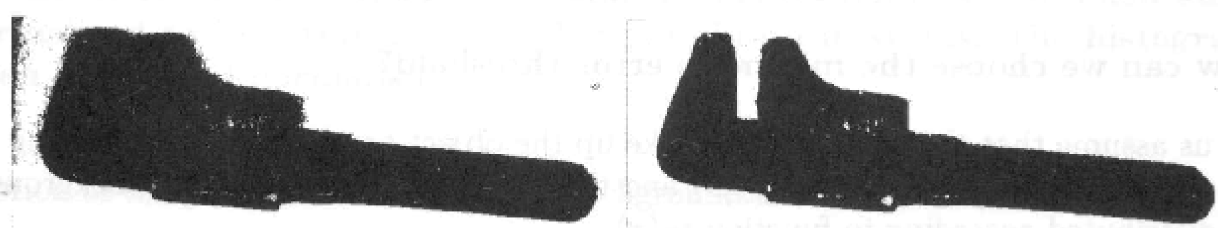

• Using prior knowledge: the P-Tile method

- This method requires knowledge about the area or size of the objects present in the image

- Let us assume that we have dark objects against a light background.

- If, for example, the objects occupy p% of the image area, an appropriate threshold can be chosen by partitioning the histogram

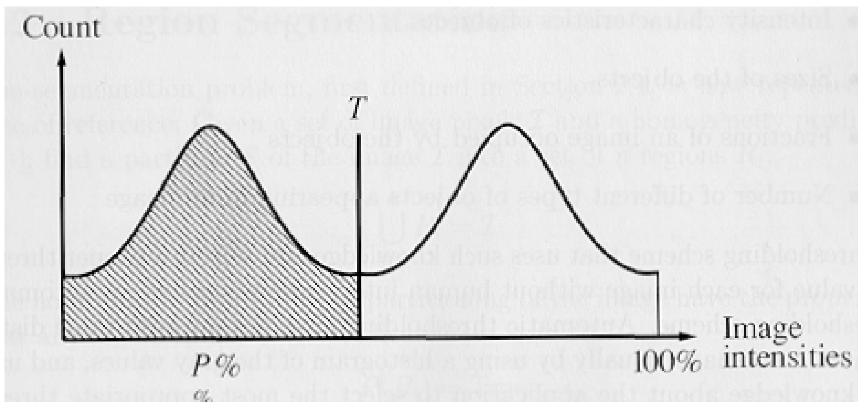

• Optimal thresholding

- Suppose that an image contains only two principal regions (e.g., object and background)

- We can minimize the number of misclassified pixels if we have some prior knowledge about the distributions of the gray level values that make up the object and the background.

- Assume that the distribution of gray-level values in each region follows a Gaussian distribution

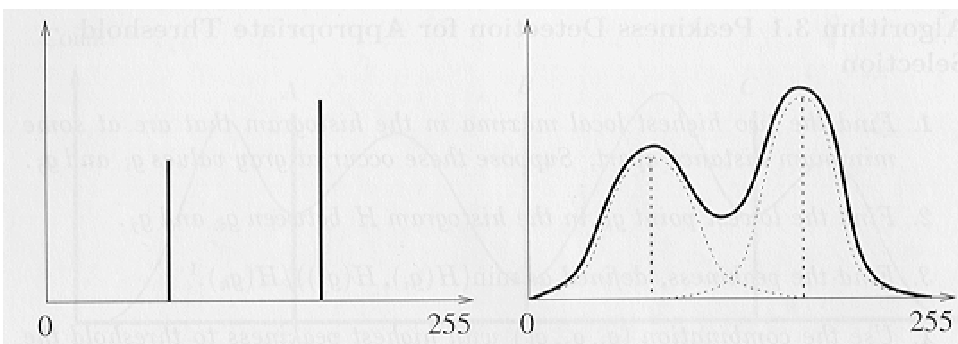

- The probability of a pixel value is then given by the following mixture:

$$P(z) = p(z/background)\ P(background) + p(z/object)\ P(object)$$

$$\text{or } P(z) = P_b \frac{1}{\sqrt{2\pi\sigma_b}} e^{-\frac{(z-\mu_b)^2}{2\sigma_b^2}} + P_o \frac{1}{\sqrt{2\pi\sigma_o}} e^{-\frac{(z-\mu_o)^2}{2\sigma_o^2}}$$

$$\text{or } P(z) = P_b p_b(z) + P_o p_o(z)$$

$p_b(z)$, $p_o(z)$, prob. distributions of background, object pixels

μ_b, μ_o: the means of the distributions

σ_b, σ_o: the standard deviations of the distributions

P_b, P_o: the a-priori probabilities of background, object pixels

- Suppose we have chosen a threshold T, what is the probability of (erroneously) classifying an object pixel as background ?

$$E_o(T) = \int_{-\infty}^{T} p_o(z)dz$$

- What is the probability of (erroneously) classifying a background pixel as object ?

$$E_b(T) = \int_{T}^{\infty} p_b(z)dz$$

- Overall probability of error

$$E(T) = P_o E_o(T) + P_b E_b(T)$$

- Minimize $E(T)$

$$\frac{dE(T)}{dT} = 0$$

- The above expression is minimized when

$$T = \frac{\mu_b + \mu_o}{2} + \frac{\sigma^2}{\mu_b - \mu_o} \; ln(P_o/P_b) \;\; (\sigma_b = \sigma_o = \sigma)$$

- If P_b=P_o or if σ=0, then

$$T = \frac{\mu_b + \mu_o}{2}$$

- How can we choose T in practice ?

(1) Find the histogram $h(z)$ of the image to be segmented

(2) Choose the parameters (μ_b, μ_o, σ_b, σ_o, P_b, P_o) such that the model $p(z) = P_b p_b(z) + P_o p_o(z)$ fits $h(z)$ satisfactorily

e.g., minimize $Error = \frac{1}{N} \sum_{i=1}^{N} (p(z_i) - h(z_i))^2$

(3) Choose T based on the formula derived above

- Drawbacks of the optimum thresholding method

  * Prior probabilities might not be known.

  * Object/Background distributions might not be known.

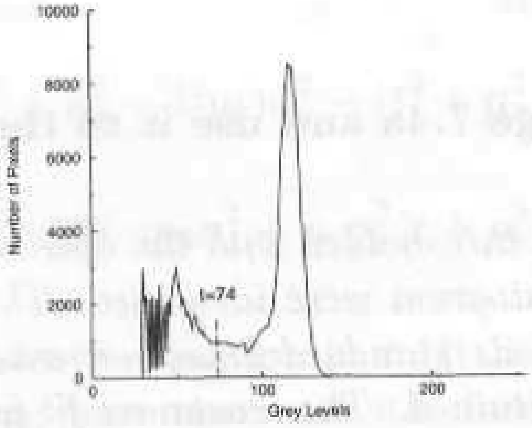

(a) Histogram with the optimal threshold marked

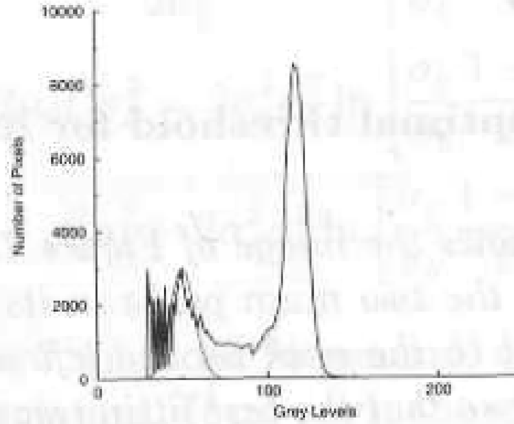

(b) Histogram with the Gaussian model for the object pixels superimposed

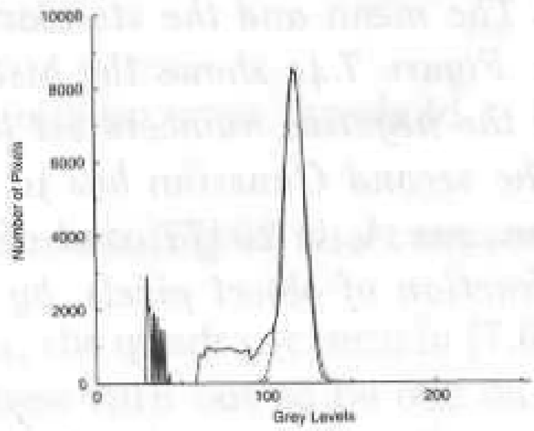

(c) Histogram after subtraction of the Gaussian used to model the object pixels with the Gaussian model for the background pixels superimposed.

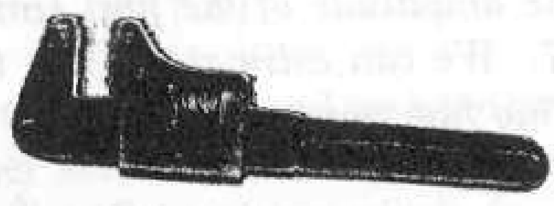

(d) Image thresholded with the optimal threshold

• **Otsu's method** (Petrou et al., pp. 278-282)

- A measure of region homogeneity is variance (i.e., regions with high homogeneity will have low variance).

- Otsu's method selects the threshold by minimizing the within-class variance of the two groups of pixels separated by the thresholding operator.

- It does not depend on modeling the probability density functions, however, it assumes a bimodal distribution of gray-level values (i.e., if the image approximately fits this constraint, it will do a good job).

• **Means and variances**

- Consider that we have an image with L gray levels and its normalized histogram (i.e., for each gray-level value i, $P(i)$ is the normalized frequency of i).

- Assuming that we have set the threshold at T, the -normalized- fraction of pixels that will be classified as background and object will be:

$$q_b(T) = \sum_{i=1}^{T} P(i), \qquad q_o(T) = \sum_{i=T+1}^{L} P(i) \quad (q_b(T) + q_o(T) = 1)$$

- The mean gray-level value of the background and the object pixels will be:

$$\mu_b(T) = \frac{\sum_{i=1}^{T} iP(i)}{\sum_{i=1}^{T} P(i)} = \frac{1}{q_b(T)} \sum_{i=1}^{T} iP(i) \qquad \mu_o(T) = \frac{\sum_{i=T+1}^{L} iP(i)}{\sum_{i=T+1}^{L} P(i)} = \frac{1}{q_o(T)} \sum_{i=T+1}^{L} iP(i)$$

- The mean gray-level value over the whole image (grand mean) is:

$$\mu = \frac{\sum_{i=1}^{L} iP(i)}{\sum_{i=1}^{L} P(i)} = \sum_{i=1}^{L} iP(i)$$

- The variance of the background and the object pixels will be:

$$\sigma_b^2(T) = \frac{\sum_{i=1}^{T}(i-\mu_b)^2 P(i)}{\sum_{i=1}^{T} P(i)} = \frac{1}{q_b(T)} \sum_{i=1}^{T}(i-\mu_b)^2 P(i)$$

$$\sigma_o^2(T) = \frac{\sum_{i=T+1}^{L}(i-\mu_o)^2 P(i)}{\sum_{i=T+1}^{L} P(i)} = \frac{1}{q_o(T)} \sum_{i=T+1}^{L}(i-\mu_o)^2 P(i)$$

- The variance of the whole image is:

$$\sigma^2 = \sum_{i=1}^{L}(i-\mu)^2 P(i)$$

• Within-class and between-class variance

- It can be shown that the variance σ can be written as follows:

$$\sigma^2 = q_b(T)\sigma_b^2(T) + q_o(T)\sigma_o^2(T) \quad + \quad q_b(T)(\mu_b(T)-\mu)^2 + q_o(T)(\mu_o(T)-\mu)^2 =$$

$$\sigma_W^2(T) + \sigma_B^2(T)$$

where $\sigma_W^2(T)$ is defined to be the within-class variance and $\sigma_B^2(T)$ is defined to be the between-class variance.

• Determining the threshold

- Since the total variance σ does not depend on T, the T minimizing σ_W^2 will be the T maximizing σ_B^2.

- Let's consider maximizing σ_B^2, we can rewrite σ_B^2 as follows:

$$\sigma_B^2 = \frac{[\mu(T) - \mu q_B(T)]^2}{q_B(T) q_o(T)}$$

where $\mu(T) = \sum_{i=1}^{T} iP(i)$

- Start from the beginning of the histogram and test each gray-level value for the possibility of being the threshold T that maximizes σ_B^2

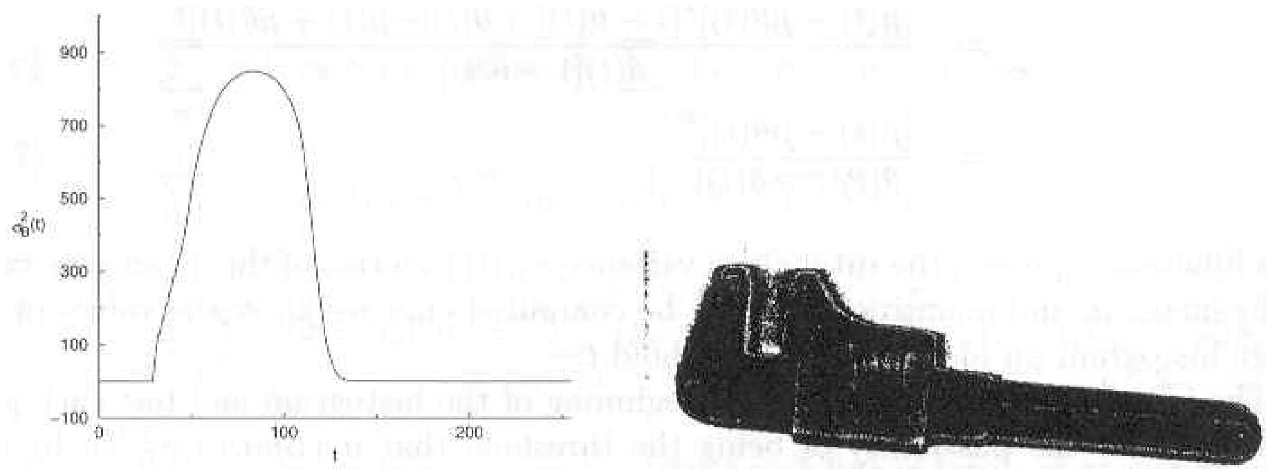

(a) $\sigma_B(t)$ *versus* t. (b) *Image thresholded with Otsu's threshold.*

• Drawbacks of the Otsu's method

- The method assumes that the histogram of the image is bimodal (i.e., two classes).

- The method breaks down when the two classes are very unequal (i.e., the classes have very different sizes).

    * In this case, σ_B^2 may have two maxima.

    * The correct maximum is not necessary the global one.

    * The selected threshold should correspond to a valley of the histogram.

- The method does not work well with variable illumination.

• Effect of illumination on segmentation

$$f(x, y) = i(x, y) r(x, y)$$

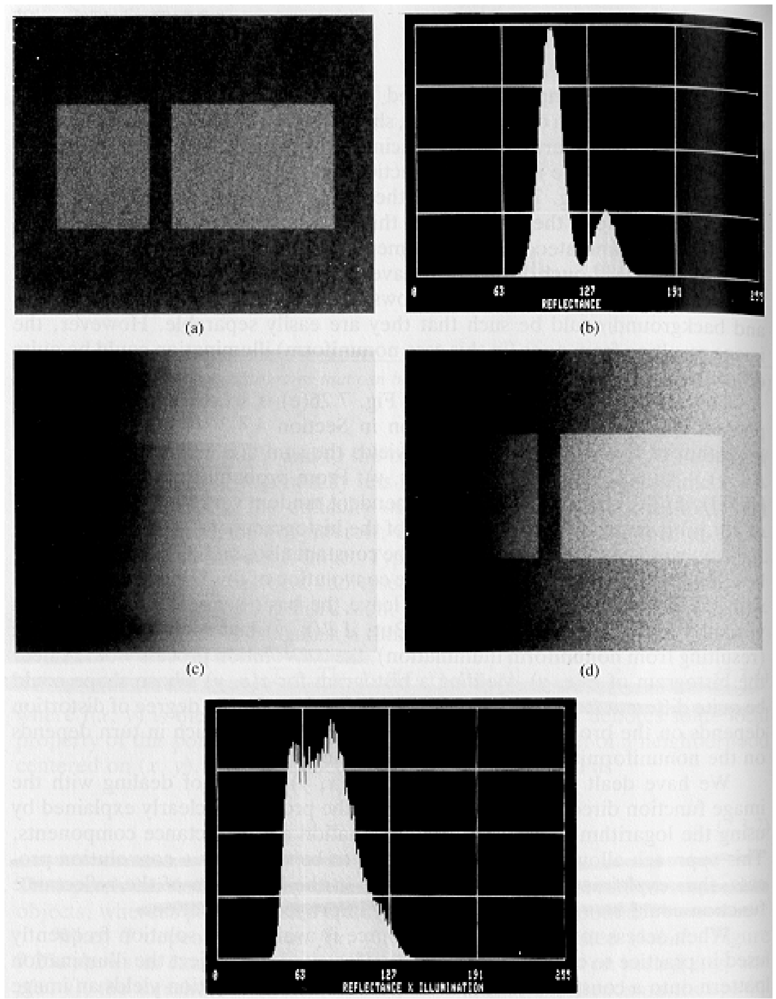

- How does illumination affect histogram ?

$$ln(f(x, y)) = ln(i(x, y)) + ln(r(x, y))$$

$$hist(ln(f(x, y))) = hist(ln(i(x, y))) + hist(ln(r(x, y)))$$

• Handling nonuniform illumination: laboratory solution

- Obtain an image of just the illumination field.

1. Suppose that $f(x, y) = i(x, y)r(x, y)$, where $i(x, y)$ is non-uniform

2. Project the illumination pattern on a surface with uniform reflectance (e.g., a white surface)

$$g(x, y) = k\ i(x, y)$$

3. Normalize $f(x, y)$:

$$h(x, y) = f(x, y)/g(x, y) = r(x, y)/k$$

4. If $r(x, y)$ can be segmented using T, then $h(x, y)$ can be segmented using T/k

• Handling nonuniform illumination: local thresholding

- A single threshold will not work well when we have uneven illumination due to shadows or due to the direction of illumination.

- The idea is to partition the image into $m\ x\ m$ subimages and then choose a threshold T_{ij} for each subimage.

- This approach might lead to subimages having simpler histogram (e.g., bimodal)

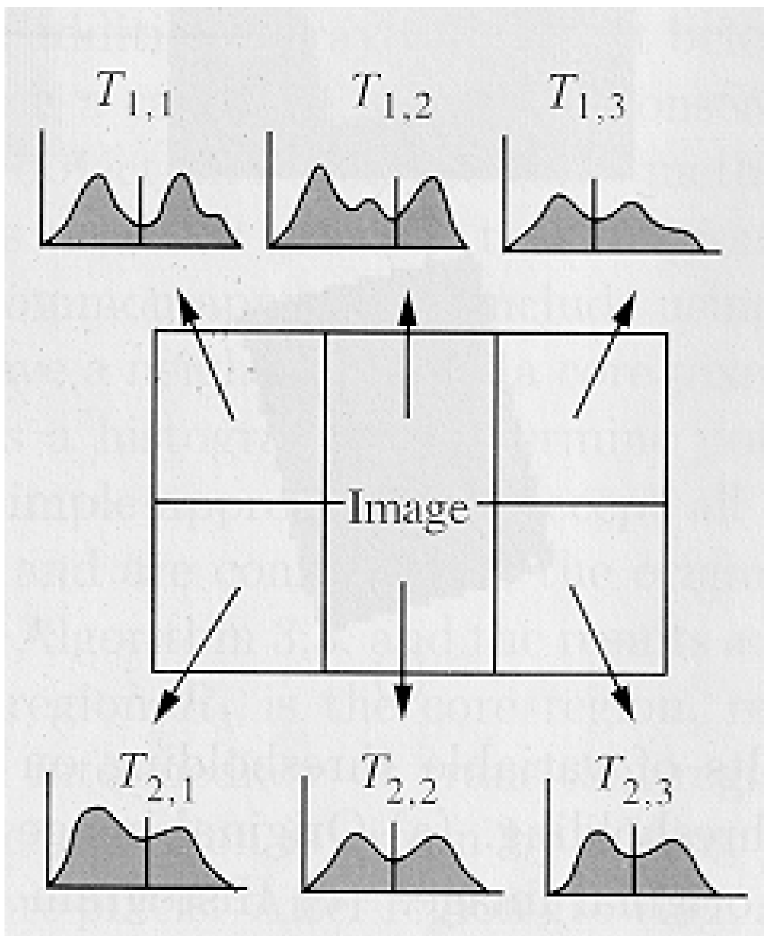

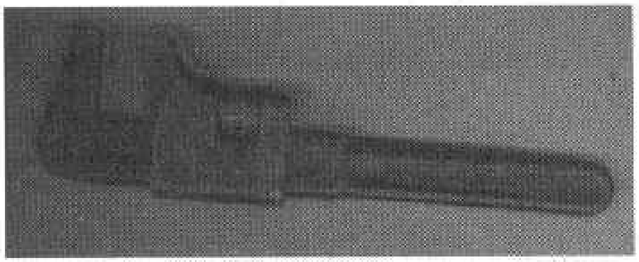

(a) Original image

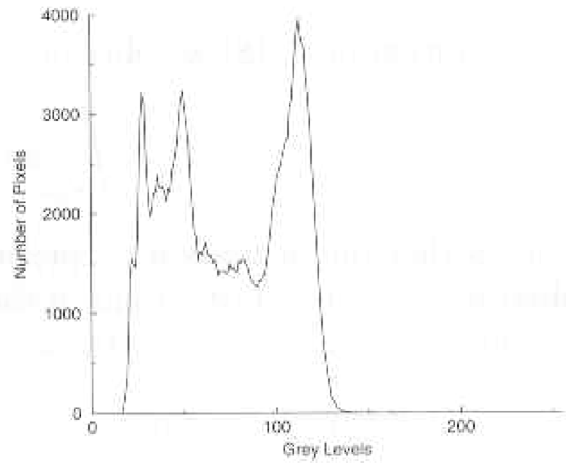

(b) Histogram

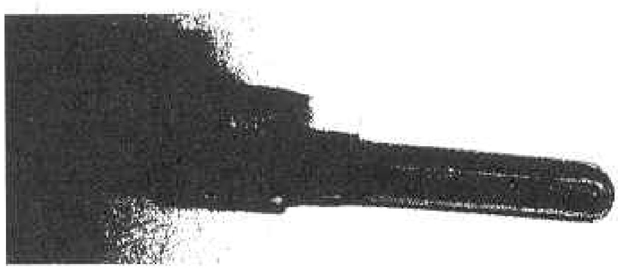

(c) Global thresholding

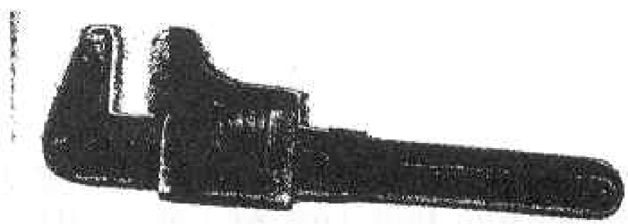

(d) Local thresholding

(using Otsu’s method and local thresholding)

• Handling nonunform illumination: variable thresholding

- In case of uneven illumination, another useful technique is to approximate the values of the image by a simple function (i.e., plane).

- Thersholding can be done relative to the plane (e.g., points above the plane will be part of the object and anything below will be part of the background).

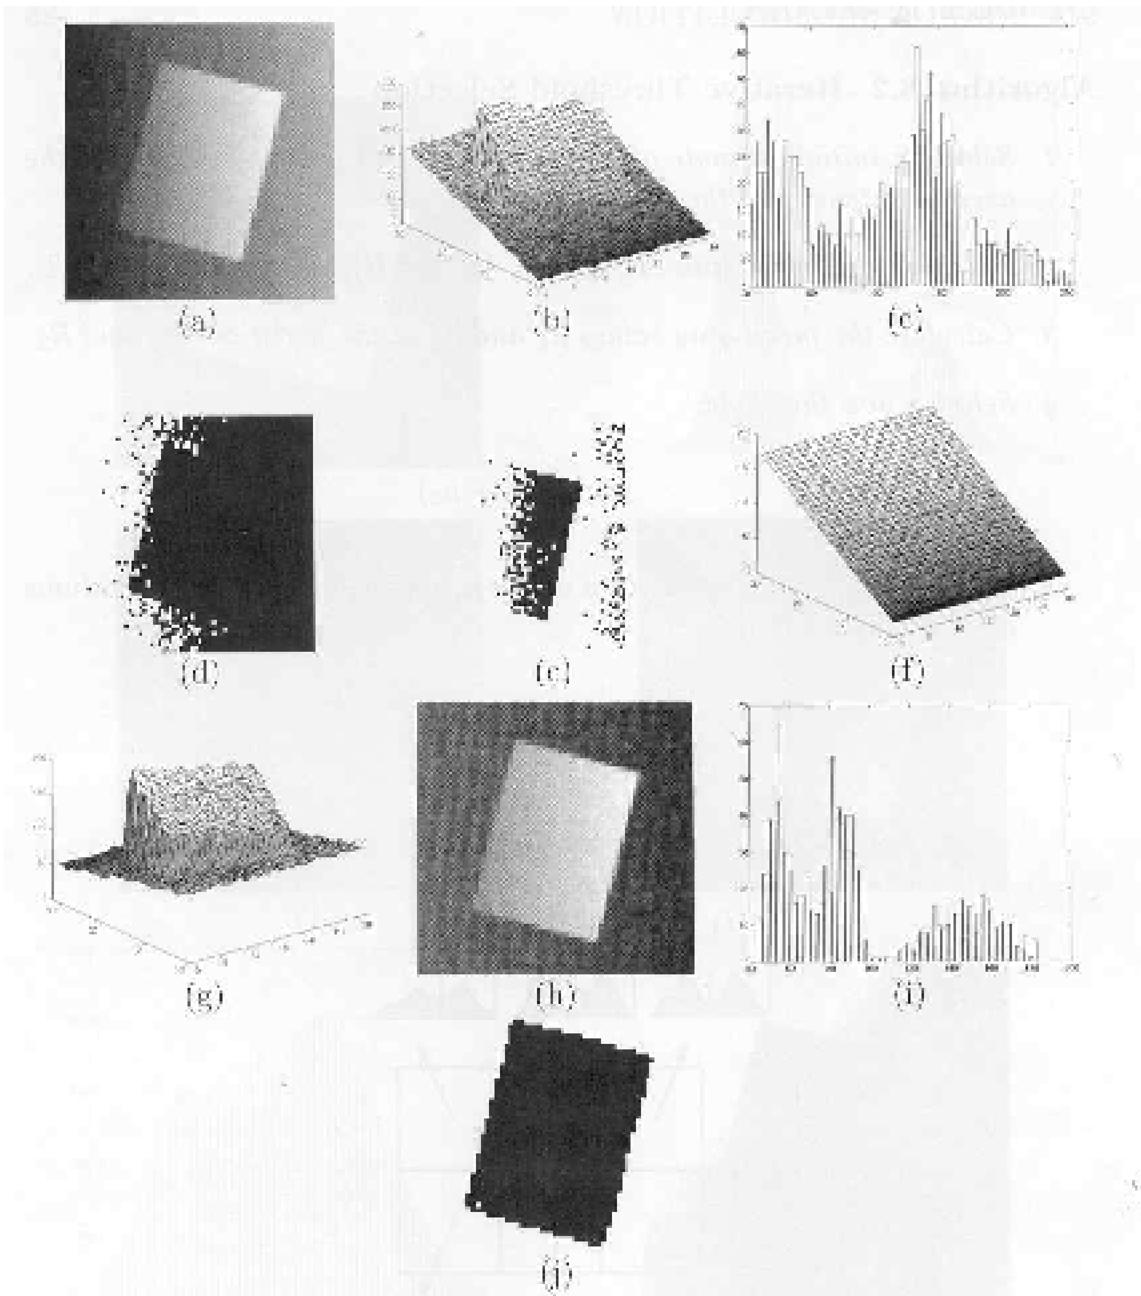

• Drawbacks of thresholding

- Pixels assigned to a single class need not form coherent regions as the spatial locations of pixels are completely ignored (*Note*: Only hysteresis thresholding considers some form of spatial proximity).

- Threshold selection is not always straightforward.