DD2423 Image Analysis and Computer Vision

DIGITAL GEOMETRY

Mårten Björkman

Computational Vision and Active Perception

School of Computer Science and Communication

November 11, 2013

Digital geometry

- Image histogram and image enhancement
- Binary images, thresholding
- Neighborhood concept
- Connectivity, connected components
- Distance measures
- Distance transforms
- Histogram equalization

Image enhancement by gray-level transformations

- Contrast enhancement
- Histogram equalization

Image histogram:

- Provides information about the contrast and overall intensity distribution.
- Simply a bar graph of pixel intensities.

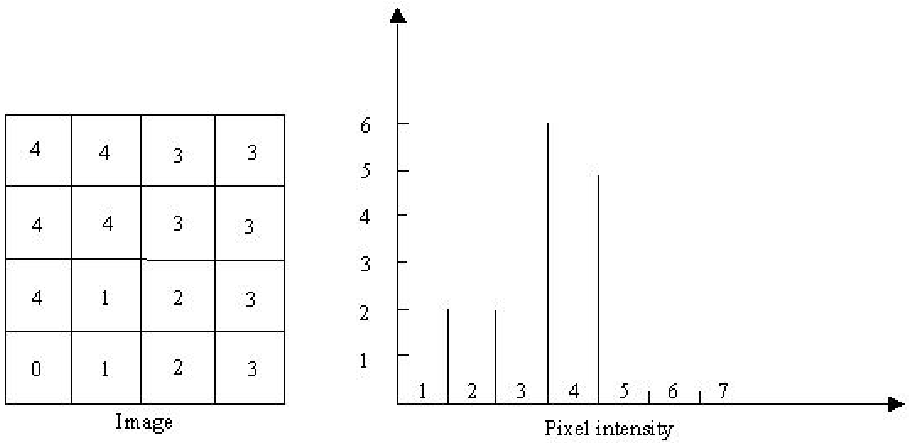

Binary images

- Images with two colors, black and white.
- Gray-level is either 0 or 1 (255)
  - Commonly referred to as ’background’ and ’foreground’.
- Typically obtained from thresholding or image segmentation.

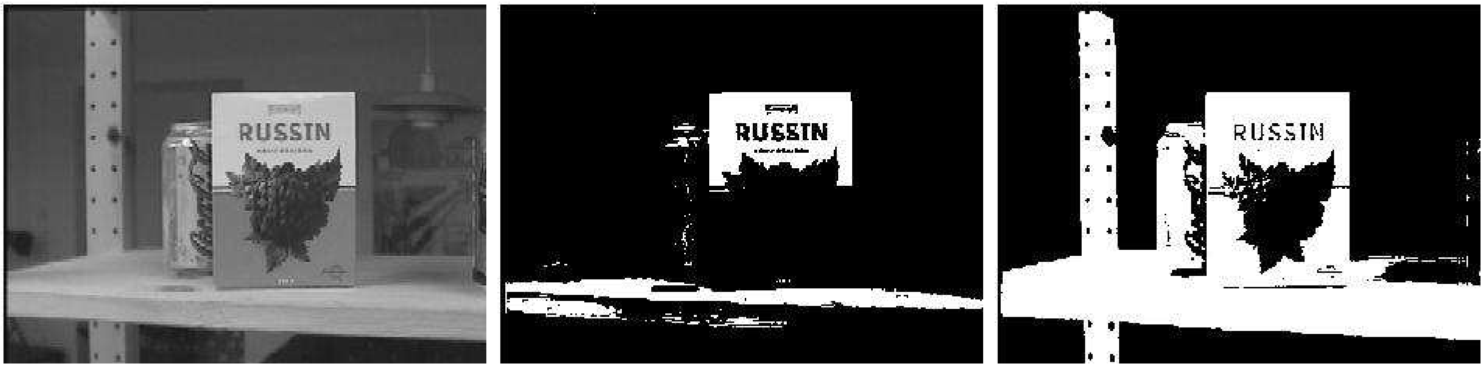

Segmentation

Simplest Image Segmentation is done by thresholding. This requires that an object has an homogenous intensity and a background with a different intensity level. Such an image can be segmented into two regions by simple thresholding:

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

Example:

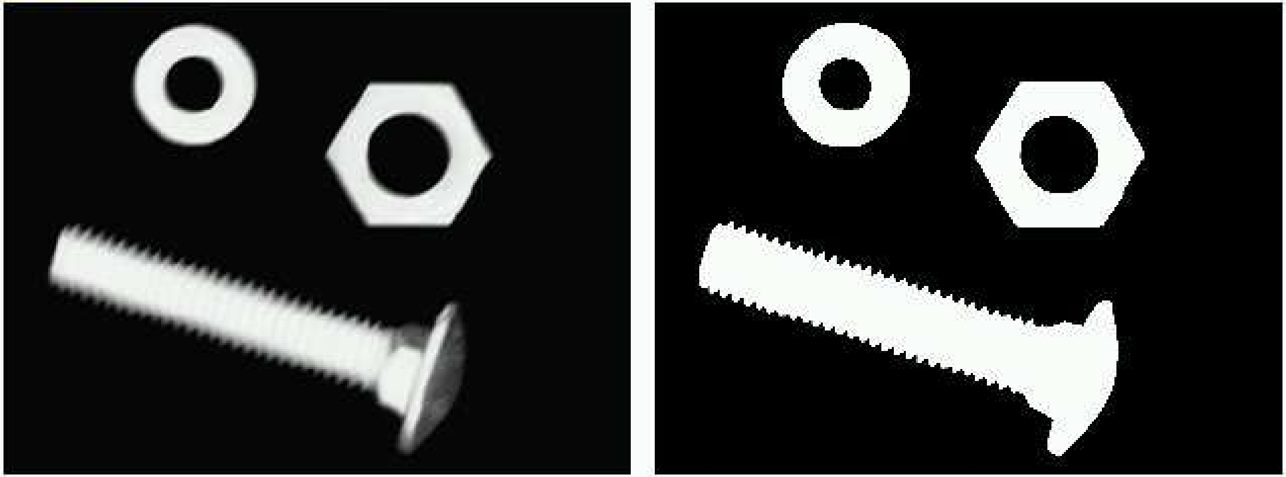

Digital geometry

Many image processing operations are based on local neighborhood operations.

Pixels are 4-neighbours if their distance is $D_4 = 1$

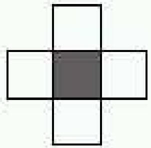

all 4-neighbours of center pixel

Pixels are 8-neighbours if their distance is $D_8 = 1$

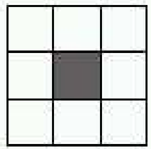

all 8-neighbours of center pixel

Connectivity

- Path: A path from p to q is a set of points $p_0 \ldots p_n$, such that each point p_i is a neighbor of p_{i-1}.
- Connectivity: p is connected to q in S, if there is a path from p to q completely in S.

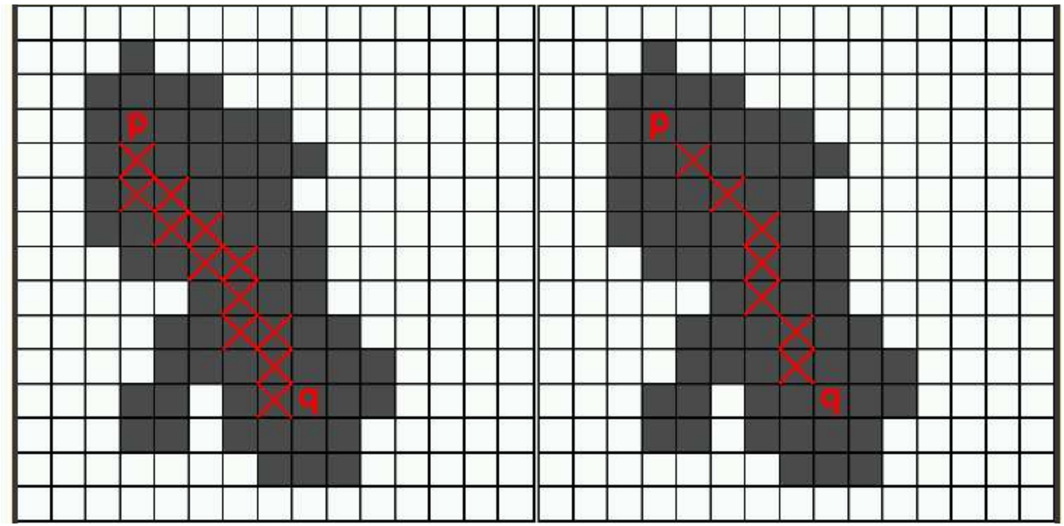

Connected components

- For every p, the set of all points q connected to p is said to be its connected component.

Recursive procedure that scans entire image:

1. for each unlabeled foreground pixel, assign it a new label L
2. assign label L to all neighboring foreground pixels
3. stop if there is no unlabeled foreground pixels

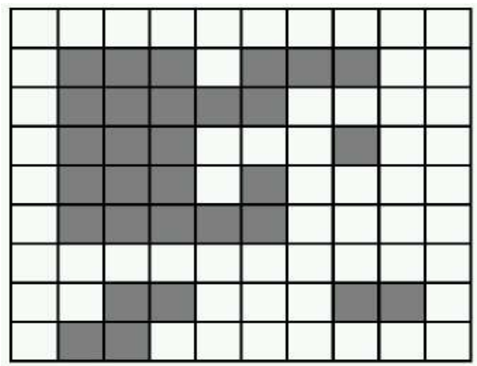

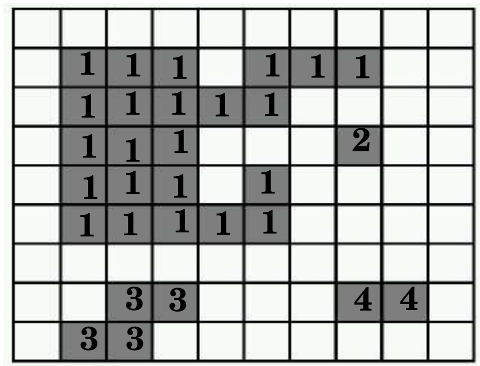

Connected component labeling

Regions (connected components) are often denoted by labels.

- statistics of regions (size, shape, gray-level statistics)
- size filtering (suppress objects of size $<$ threshold)

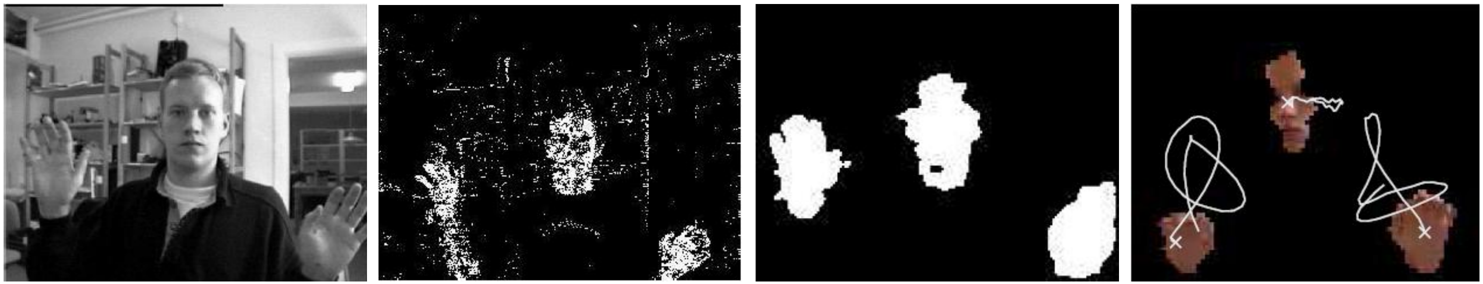

Duality of 4-connectivity and 8-connectivity

Outer boundary: set of background points with a neighbor on the object.

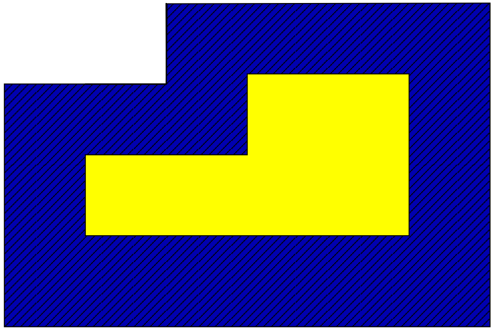

(left) based on 8-connectivity

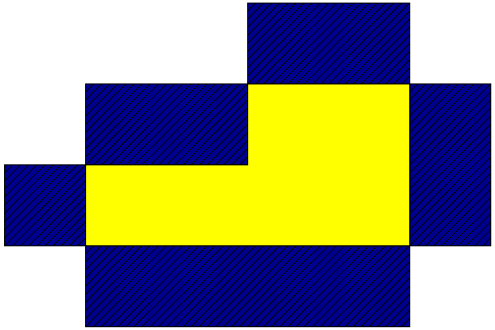

(right) based on 4-connectivity

- Jordan curve theorem (continuous case):
  Each closed curve divides plane into one region inside and one region outside.
- Note: Many region based methods, only store the boundary.

Duality of 4-connectivity and 8-connectivity

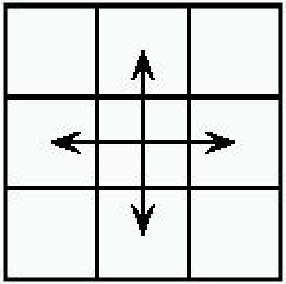

Figure 2.1: 4-connected neighbors.

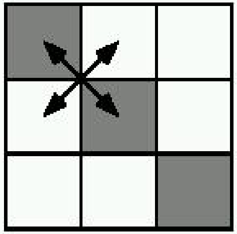

Figure 2.2: Paradox of 4-connected neighbors.

If you separate two 4-connected regions, the boundary between them needs to be 8-connected, and if you separate two 8-connected regions, the boundary between them needs to be 4-connected.

But, what if the black pixels are (to be considered as) the background, while the white are the foreground?

Duality of 4-connectivity and 8-connectivity

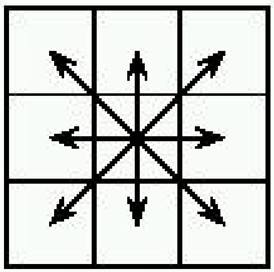

Figure 2.3: 8-connected neighbors.

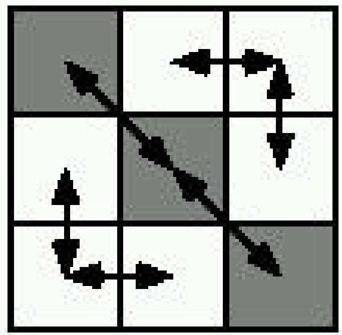

The usual solution is to use 4-connectivity for the foreground with 8-connectivity for the background or to use 8-connectivity for the foreground with 4-connectivity for the background.

Duality of 4-connectivity and 8-connectivity

- If 4(8)-connectivity used for foreground

  $\Rightarrow$ 8(4)-connectivity used for background

- Hexagonal grid - same connectivity concept for foreground and background

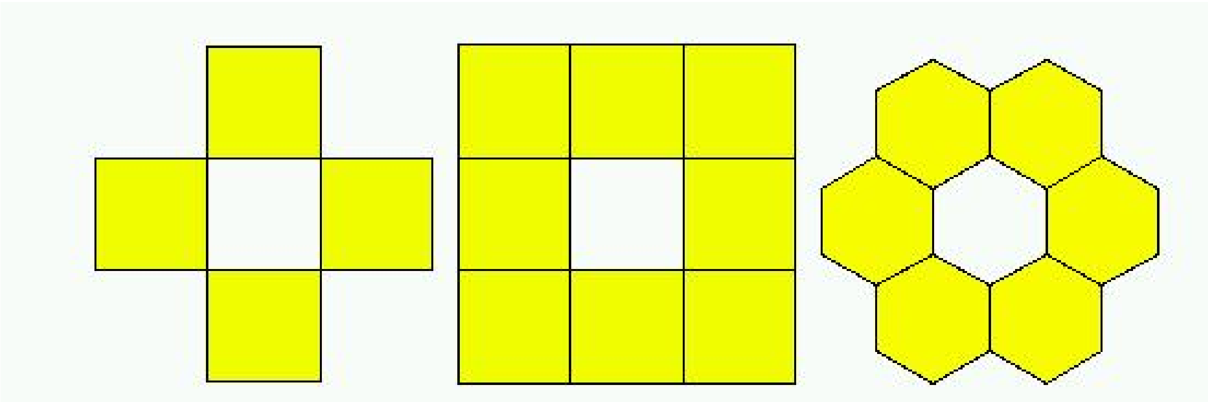

Distance measures

How to define distance between two points p and q?

Common distance measures:

- Euclidean distance $d(p,q) = \sqrt{(x-u)^2 + (y-v)^2}$
- City block distance $d(p,q) = |x-u| + |y-v|$
- Chessboard distance $d(p,q) = \max(|x-u|, |y-v|)$

All three measure satisfy metric axioms

- $d(p,q) \geq 0$
- $d(p,q) = d(q,p)$
- $d(p,r) \leq d(p,q) + d(q,r)$

Distance measures

Euclidean distance

$\sqrt{8}$	$\sqrt{5}$	2	$\sqrt{5}$	$\sqrt{8}$
$\sqrt{5}$	$\sqrt{2}$	1	$\sqrt{2}$	$\sqrt{5}$
2	1	0	1	2
$\sqrt{5}$	$\sqrt{2}$	1	$\sqrt{2}$	$\sqrt{5}$
$\sqrt{8}$	$\sqrt{5}$	2	$\sqrt{5}$	$\sqrt{8}$

City block distance

4	3	2	3	2
3	2	1	2	3
2	1	0	1	2
3	2	1	2	3
4	3	2	3	4

Chessboard distance

2	2	2	2	2
2	1	1	1	2
2	1	0	1	2
2	1	1	1	2
2	2	2	2	2

Distance transform

- The result is an image that shows the distance to the closest boundary from each point
- Useful for shape description, matching, skeletonization, etc

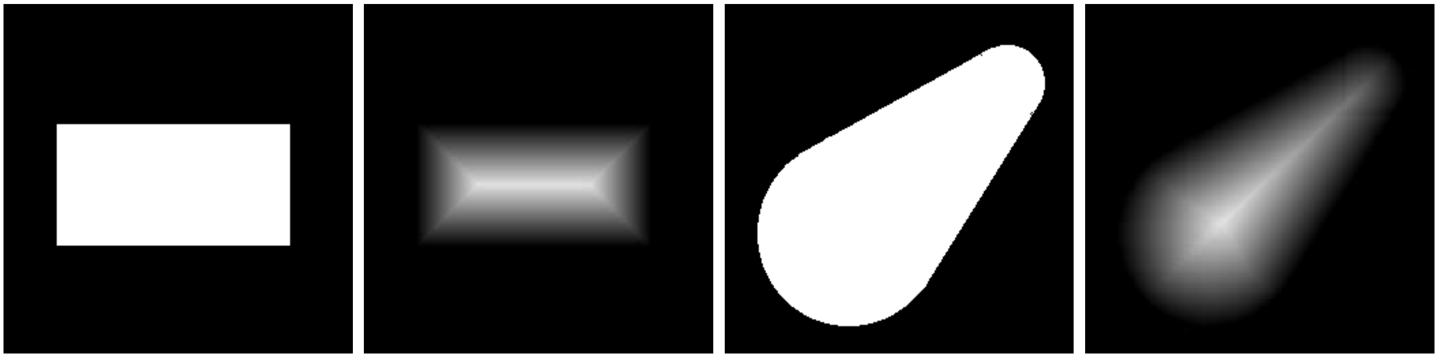

Distance transform for matching shapes

- Create distance transform from model shape S_{model} represented by edges.
- Extract new shape S_{image} from an image.
- Sum values in distance transform over edge points from S_{image}.
- Iteratively transform S_{image} until sum is mimimized.

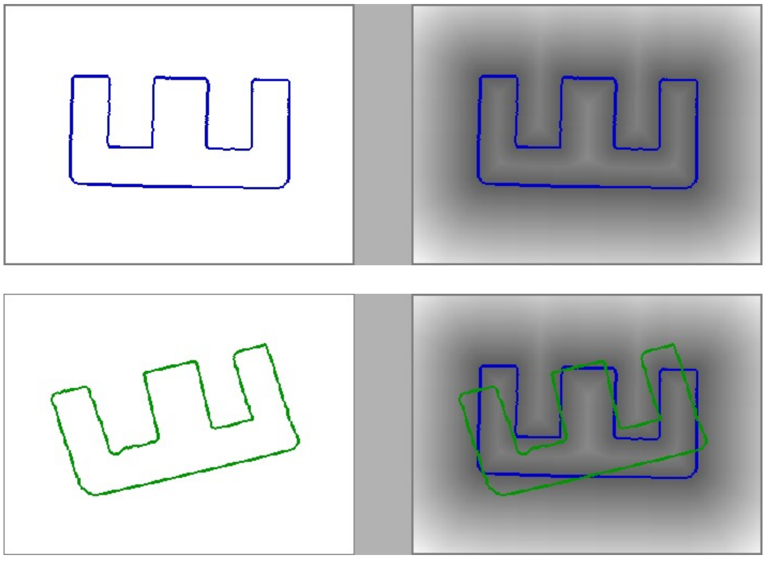

Image enhancement

- Goal: Improve the subjective quality of the image.
- Examples:

  - Contrast enhancement

  - Noise suppression - smoothing

  - Sharpening

  - Feature enhancement
- Assumption:

  - no degradation model, otherwise its called restoration.

Histogram

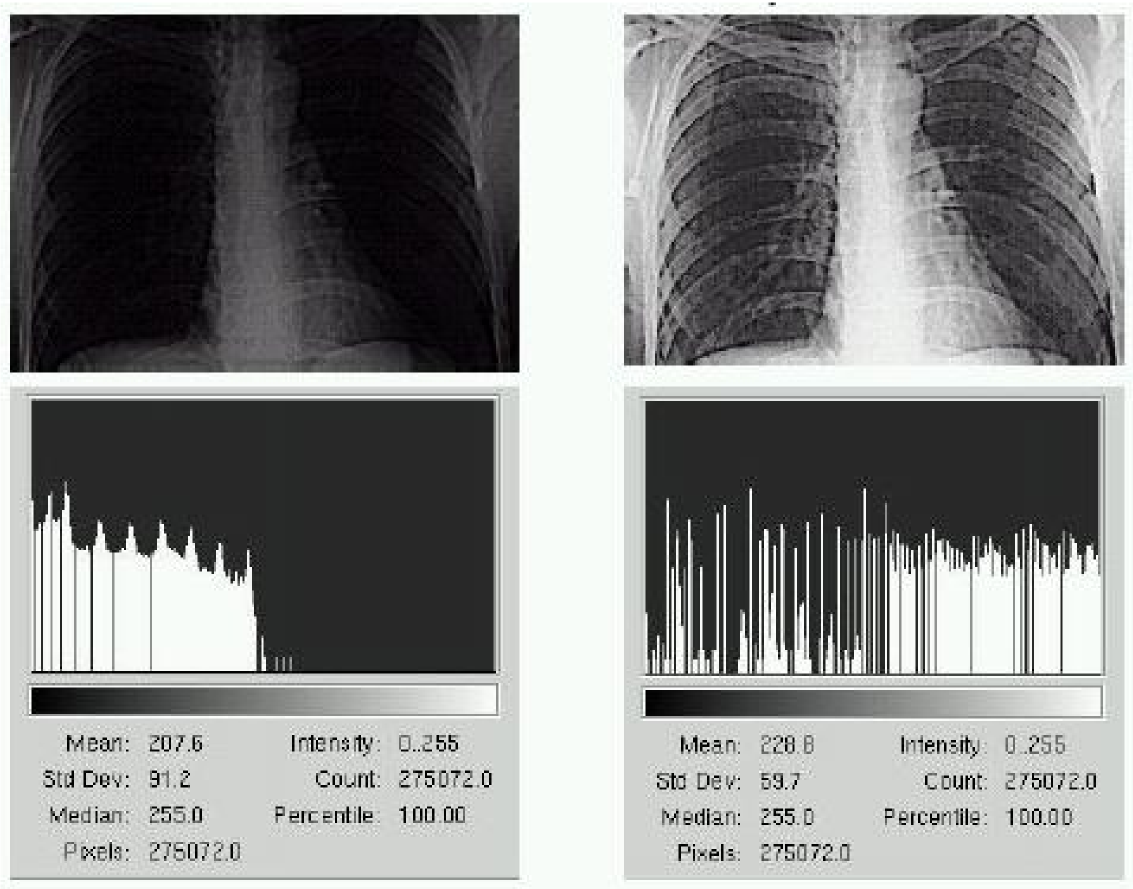
Mean: 207.6
Std Dev: 91.2
Median: 255.0
Pixels: 275072.0
Intensity: 0..255
Count: 275072.0
Percentile: 100.00
Mean: 228.8
Std Dev: 59.7
Median: 255.0
Pixels: 275072.0
Intensity: 0..255
Count: 275072.0
Percentile: 100.00

Modification of gray-levels

- Gray-level correction (position dependent): Compensate for spatially varying illumination or exposure.

$$f'(x,y) = h(x,y)f(x,y)$$

- Temporal smoothing: average multiple exposures of static scene.
- Gray-level transformations (position independent):

$$s = T(r)$$

where s and r and intensities after and before, and T may be

  – piecewise linear, negative, logarithm or power-law transformations.

Pixel Processing

Eg: take 1 pixel case

$s = \mathrm{T}(r)$

s- output gray level, r-input gray level

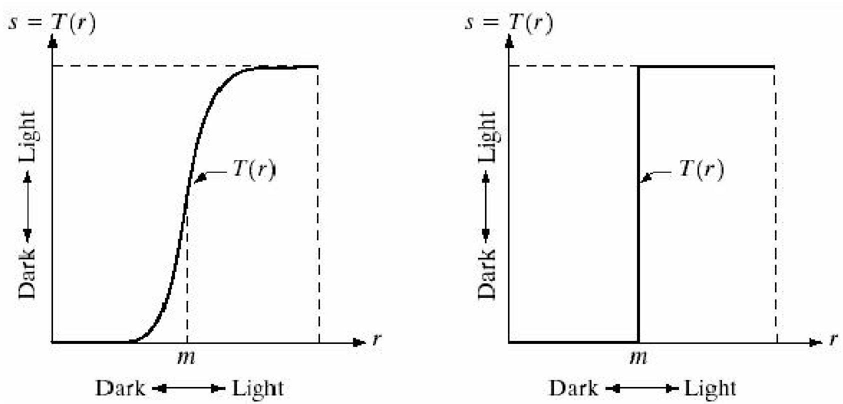

a b

FIGURE 3.2 Gray-level transformation functions for contrast enhancement.

These are point to point intensity transforms

Look-Up Tables (LUT)

Often implemented with LUTs (256 entries), at least for complex functions.

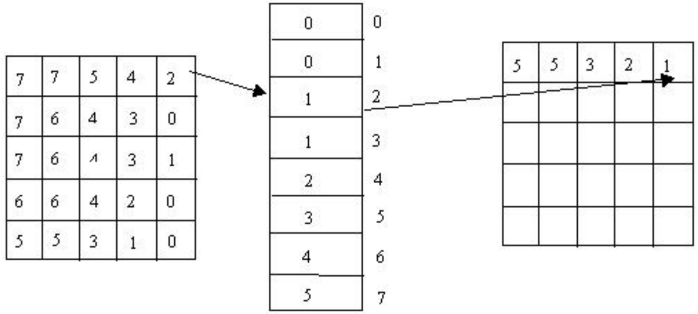

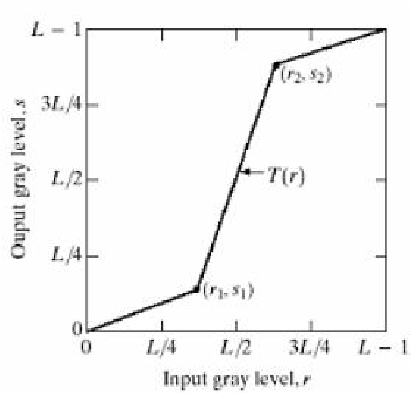

Image Negative

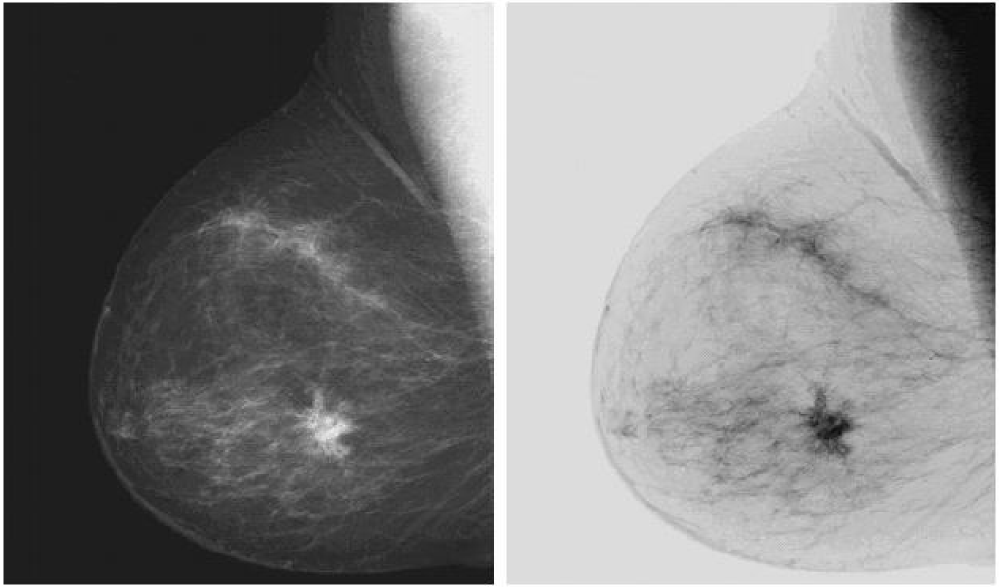

a b

FIGURE 3.4 (a) Original digital mammogram. (b) Negative image obtained using the negative transformation in Eq. (3.2-1). (Courtesy of G.E. Medical Systems.)

Image Negatives

(b) is simply the negative of (a), ie

$$s = L - 1 - r$$

where $L = 2^k$

L: # of gray levels

k: # of bits

Log transformations

Useful for compressing large dynamic range and make details visible.

$$s = c\log(1+r)$$

Example: Fourier spectrum

$$0 \rightarrow 1.5x10^6 \quad \text{to} \quad 0 \rightarrow 6.2$$

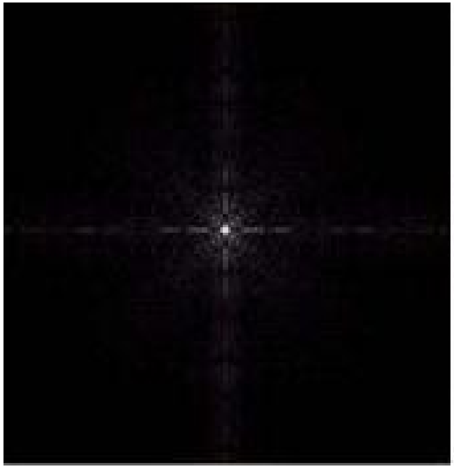

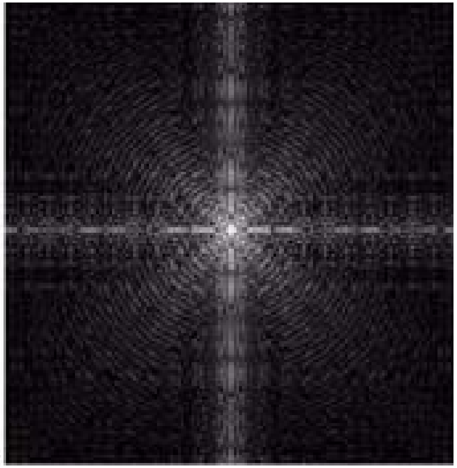

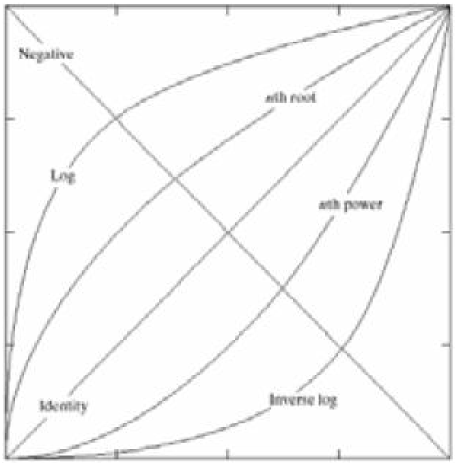

Power-law transformations

A variety of devices used for image capture, printing, and display respond according to a power law.

$$s = c\, r^{\gamma} \quad \text{or} \quad s = c\,(r + \varepsilon)^{\gamma}$$

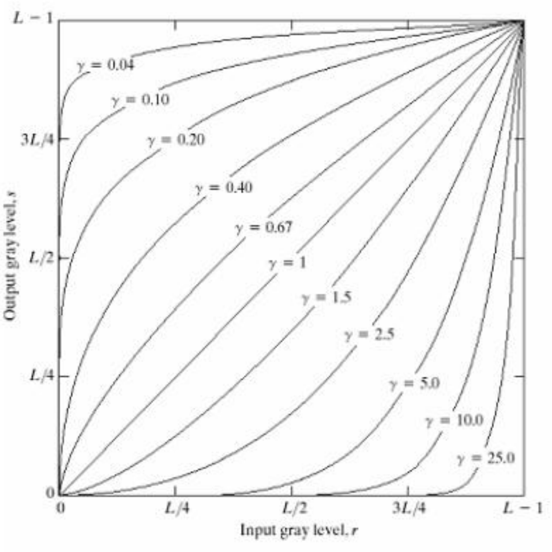

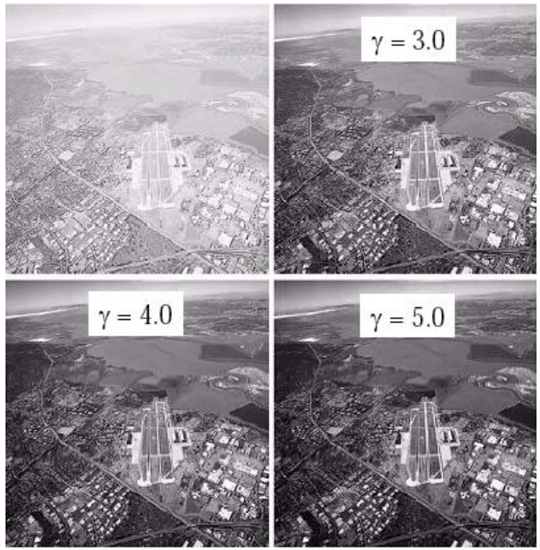

Histogram Stretching

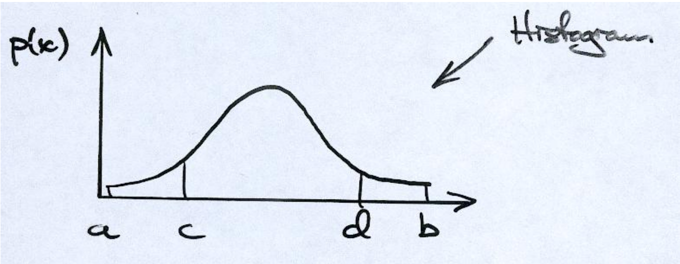

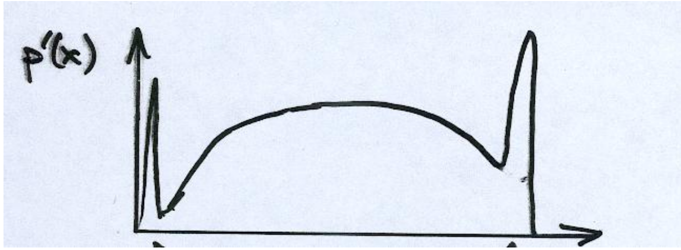

Increase contrast by letting the interval $[c, d]$ cover the entire gray-level range.
Note: Information loss in $[a, c]$ and $[d, b]$.

Full-Scale Histogram Stretch

$(r_1, s_1) = (r_{min}, 0)$

$(r_2, s_2) = (r_{max}, L\text{-}1)$

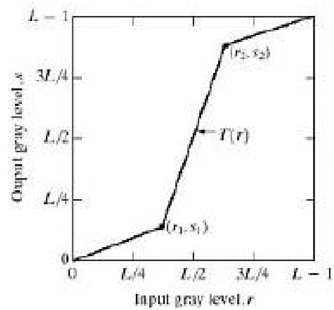

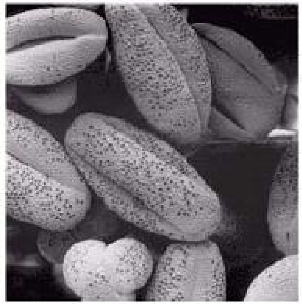

Gray-level transformations

Common requirements on transformation function $s = T(r)$:

$T(r_{min}) = r_{min}$ (or opposite) - fills up entire range of gray-levels

$T(r_{max}) = r_{max}$

T monotonic $\Rightarrow$ T invertible (no loss of information)

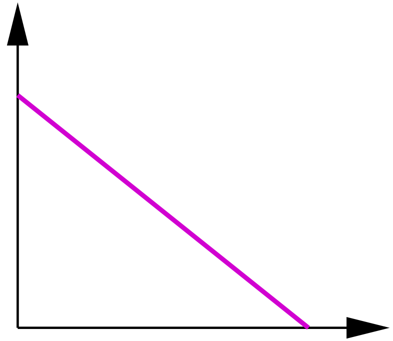

contrast reversal

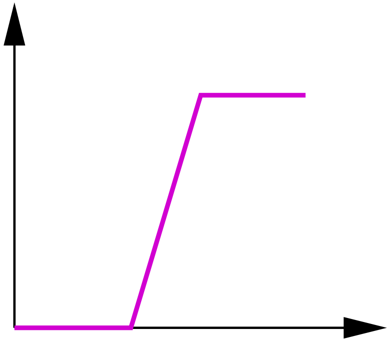

hard stretching

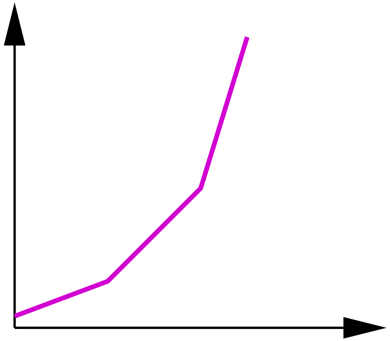

stepwise linear

Common special cases

Gray-level slicing

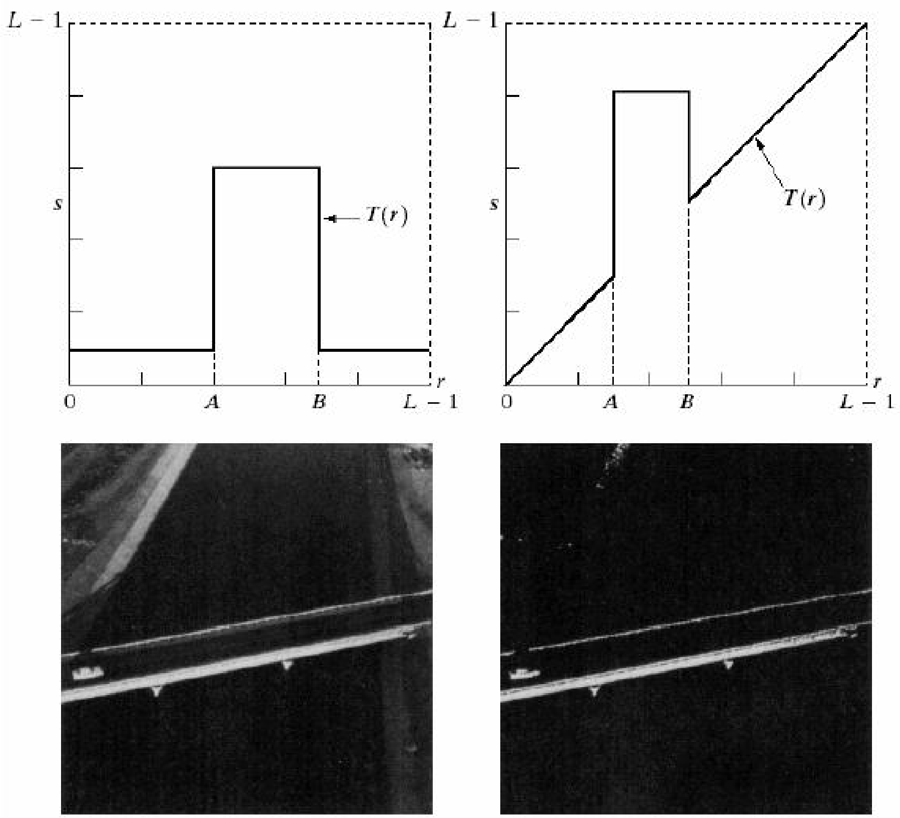

a b
c d

FIGURE 3.11
(a) This transformation highlights range [A, B] of gray levels and reduces all others to a constant level.
(b) This transformation highlights range [A, B] but preserves all other levels.
(c) An image.
(d) Result of using the transformation in (a).

Eg: Satellite imagery, X-ray imaging

Four images and their histograms

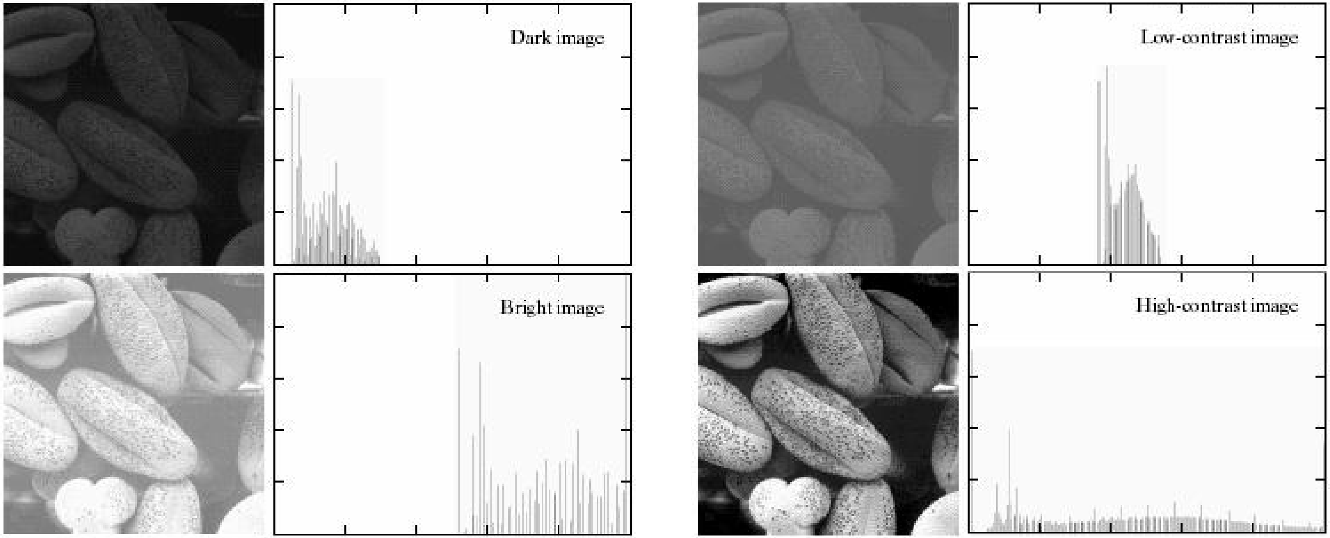

Histogram equalization

- Idea: Redistribute gray-levels as evenly as possible - this would correspond to a brightness distribution where all values are equally probable.
- Assume gray levels are continuous (not quantized) and have been normalized to lie between 0 and 1.
- Find transformation T that maps gray values r in the input image to gray values $s = T(r)$ in the transformed image.

Histogram equalization (continuous case)

We are looking for a transformation $s = T(r)$ such that the distribution $p_S(s)$ of pixel values is uniform, given a distribution from an image $p_R(r)$.

Known from probability theory:

$$p_S(s) = \left[p_R(r) \, \frac{dr}{ds} \right]_{r=T^{-1}(s)}$$

Let us define $T(r)$ as

$$s = T(r) = \int_0^r p_R(w) dw \Rightarrow \frac{ds}{dr} = p_R(r)$$

Then it follows that

$$p_S(s) = \left[p_R(r) \, \frac{dr}{ds} \right]_{r=T^{-1}(s)} = \left[p_R(r) \, \frac{1}{p_R(r)} \right]_{r=T^{-1}(s)} = 1$$

Exercise

Assume you have an image with a histogram of grey-level values given by the distribution $p_R(r) = \frac{3}{5}(4r - 4r^2 + 1),\ r \in [0,1]$. Find a transformation $s = T(r)$, such that the histogram after the transformation becomes $p_S(s) = 1$, $s \in [0,1]$.

Exercise

Assume you have an image with a histogram of grey-level values given by the distribution $p_R(r) = \frac{3}{5}(4r - 4r^2 + 1),\ r \in [0,1]$. Find a transformation $s = T(r)$, such that the histogram after the transformation becomes $p_S(s) = 1$, $s \in [0,1]$.

Answer: The transformation can be determined by computing the integral of $p_R(r)$, that is

$$s = T(r) = \int_0^r p_R(x)dx = \frac{3}{5}\left[\ \frac{4}{2}x^2 - \frac{4}{3}x^3 + x\ \right]_0^r = \frac{6}{5}r^2 - \frac{4}{5}r^3 + \frac{3}{5}r.$$

Since $T(1) = 1$, which is a condition for $p_R(r)$ to be a distribution, we don't need to normalize the transformation. Its derivative, which is the same as $p_R(r)$, determines whether the new histogram is stretched or compressed.

Histogram equalization (discrete case)

1. Compute current histogram: count each distinct pixel value in the image.
2. Store cumulative sum of all the histogram values and normalize them by multiplying each element by (maximum-pixel-value/number of pixels)

$$s_k = T(r_k) = \sum_{i=0}^{k} p_r(r_i) = \sum_{i=0}^{k} \frac{n_i}{N}, \quad 0 \leq r_k, s_k \leq 1, \;\; k = 0, 1, 2, \ldots, 255$$

3. Use LUT from step 2 to transform the input image.

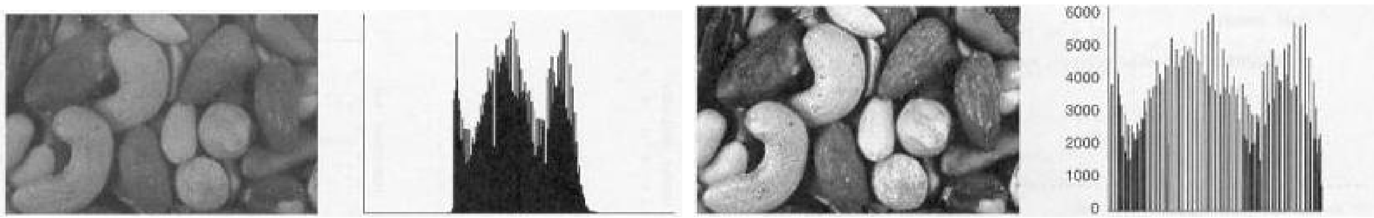

Note! Values of s_k are scaled up by 255 and rounded to the nearest integer so that the output values of this transformation will range from 0 to 255. The discretization and rounding to the nearest integer will mean that the transformed image will not have a perfectly uniform histogram.

Histogram equalization

A Worked Out Example

The table given below shows the grey level distribution of a hypothetical 32 x 32 image with 8 grey levels.

Table showing grey level distribution of input image

k	r_k	n_k	n_k / N	s_k
0	0.00	122	0.12	0.12
1	0.14	21	0.02	0.14
2	0.29	21	0.02	0.16
3	0.43	256	0.25	0.41
4	0.57	102	0.10	0.51
5	0.71	11	0.01	0.52
6	0.86	307	0.30	0.82
7	1.00	184	0.18	1.00
Total	-	N=1024	1.00	-

The following table shows the distribution of new grey levels obtained by the equalisation transformation from r_k to s_k.

Table showing grey level distribution of output image

k	s_k	n_k	n_k / N
0	0.00	0	0
1	0.14	164	0.16
2	0.29	0	0
3	0.43	256	0.25
4	0.57	113	0.11
5	0.71	0	0
6	0.86	307	0.30
7	1.00	184	0.18
Total	-	1024	1.00

Four images after histogram equalization

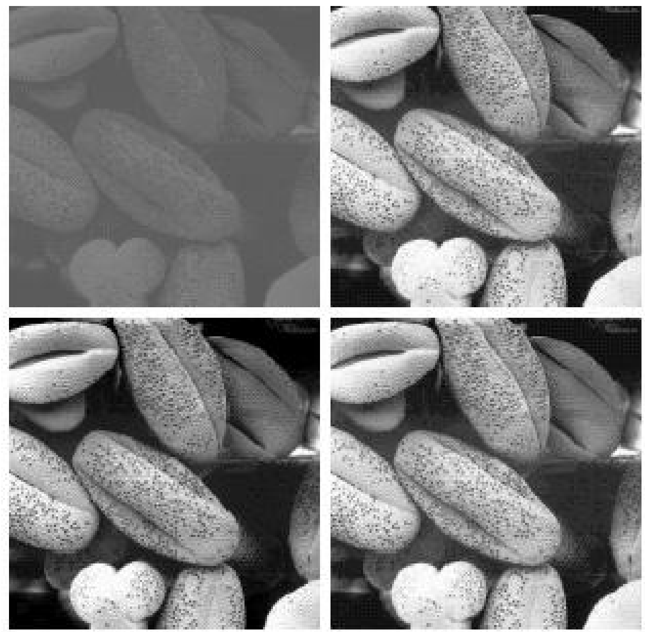

FIGURE 3.18 Transformation functions (1) through (4) were obtained from the histograms of the images in Fig.3.17(a), using Eq. (3.3-8).

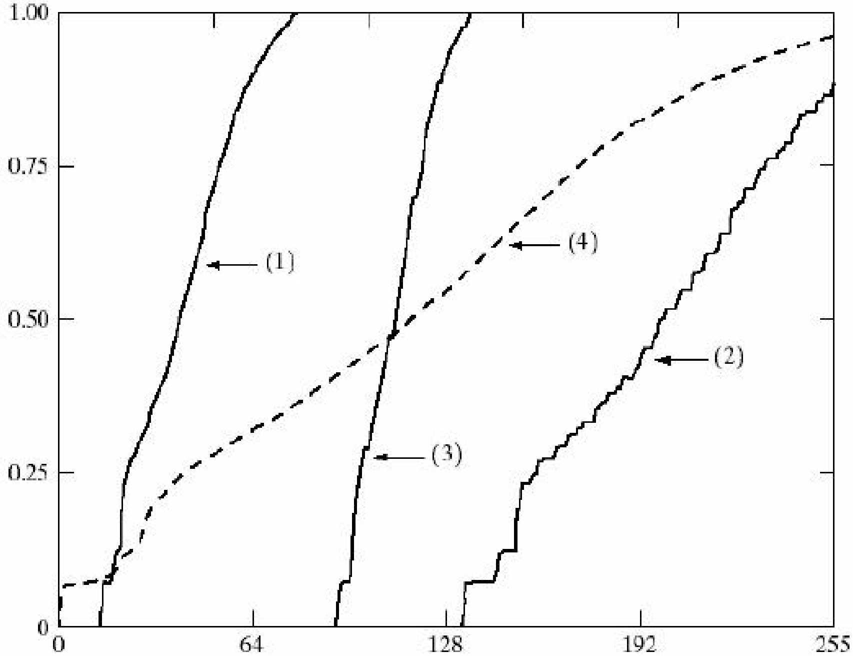

Transformation functions for previous image

Examples of transformation functions

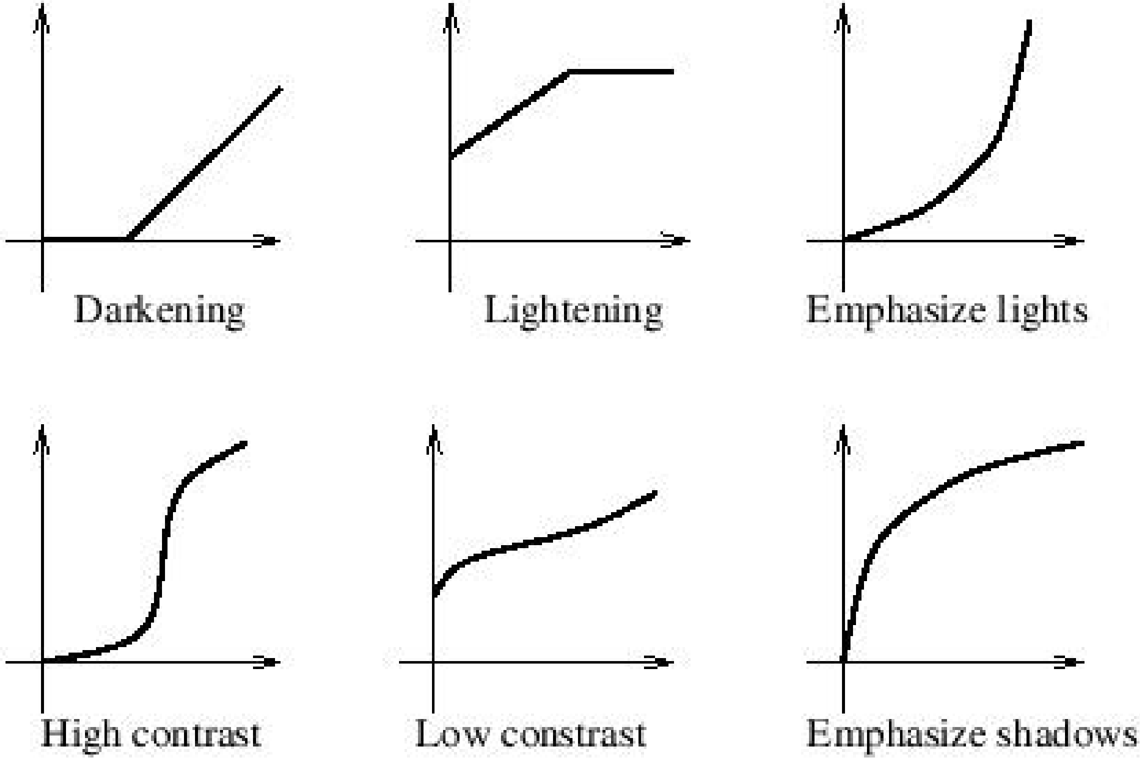

What about colours?

- The colors we see is a combination of reflectance and illumination
- Simple model: Measured colour $(R,G,B) = (I_r R_r, I_g R_g, I_b R_b)$, where (R_r, R_g, R_b) is reflectance and (I_r, I_g, I_b) illumination.
- Thus is the illumination is coloured, the measured colour changes.

What about colours?

- Grey-scale image: $I = (R+G+B)/3$ (left).
  - Can be useful for image analysis.
- Color component: $(R_n, G_n, B_n) = (R, G, B)/I$ (right).
  - Removes illumination, if illumination is assumed white, $I_r = I_g = I_b$.
  - Very difficult to use directly due to high noise level.

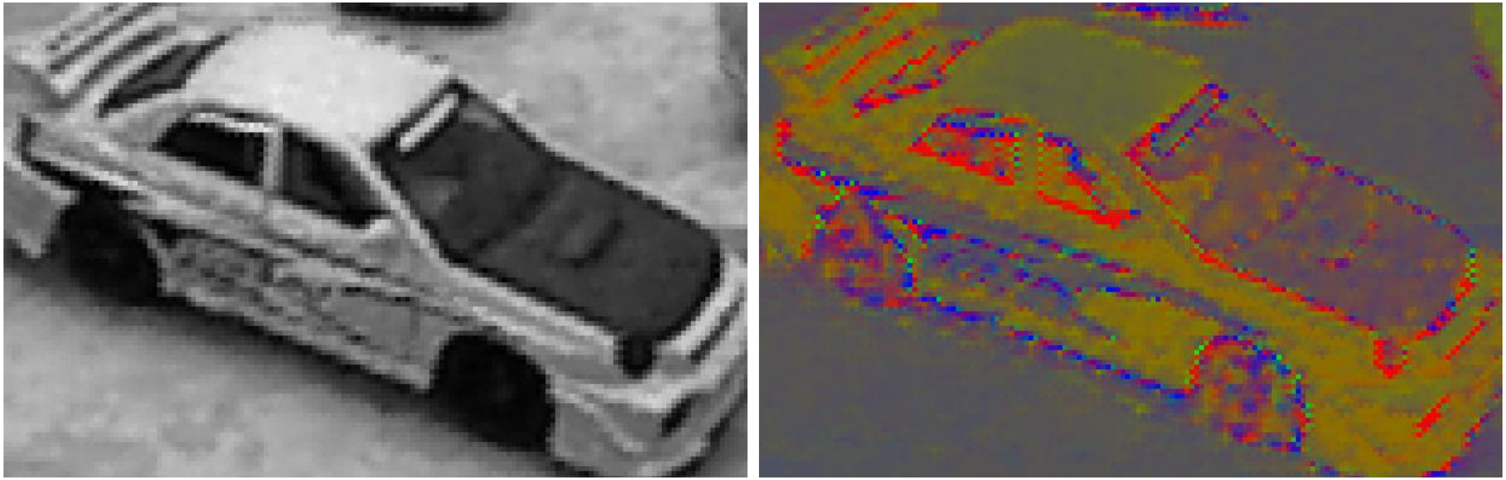

Summary of good questions

- What is a 4-neighbour and how is related to connectiveness?
- What does the duality of 4-connectivity and 8-connectivity mean?
- What kind of distance measures exist?
- Why would you like to do image enhancement?
- Mention a typical grey-level transformation. When would you use it?
- What do histogram stretching and compression mean?
- What are the principles of histogram equalization?
- Why is most image analysis done using grey-level images?

Readings

- Gonzalez and Woods: Chapter 3.2 - 3.3
- Szeliski: Chapters 3.1, 3.3.3-3.3.4