5 – IMAGE ENHANCEMENT

Image enhancement refers to accentuation or sharpening of image features such as edges, boundaries, or contrast to make a graphic display more useful for display and analysis. The enhancement process does not increase the information content in the image data, but it does increase the dynamic range of the chosen features.

Point Operations

In point operations, a given gray level r is mapped into a gray level s according to a transformation

$$s = T(r) \tag{1}$$

Some fairly simple, yet powerful, processing can be formulated with gray-level transformations.

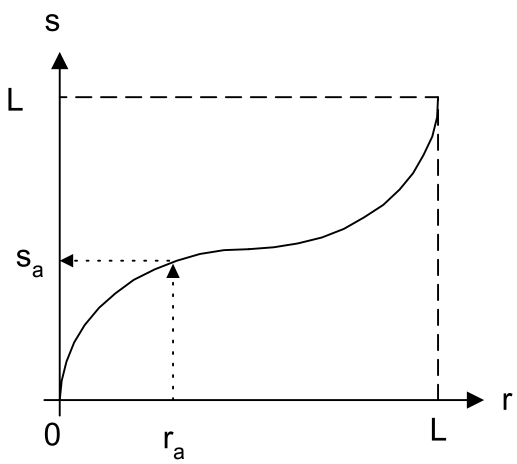

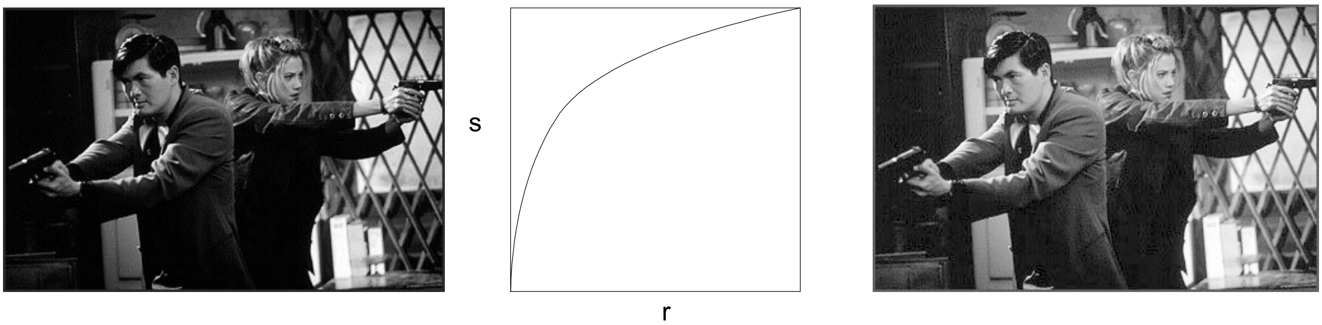
s
r

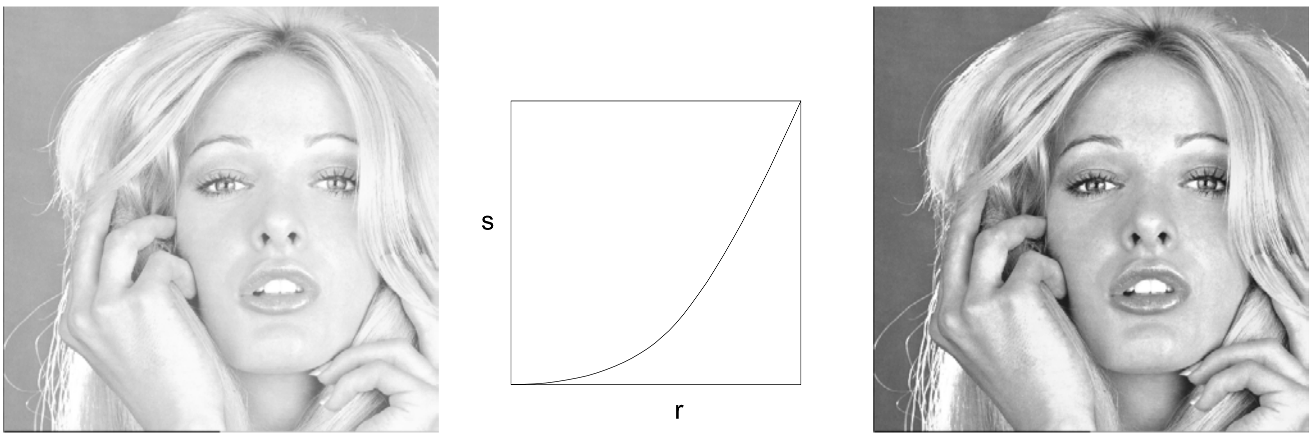
s
r

Contrast Stretching

Low-contrast images can result from poor or non-uniform lighting conditions, or due to nonlinearity or small dynamic range of the imaging sensor. A typical contrast stetching transformation can be expressed as

$$s = \begin{cases} \alpha r & 0 \le r < a \\ \beta(r-a) + s_a & a \le r < b \\ \gamma(r-b) + s_b & b \le r < L \end{cases} \tag{2}$$

The slope of the transformation is chosen greater than unity in the region of stretch. The parameters a and b can be selected by examining the histogram. For example, the gray scale intervals where pixels occur most frequently would be stretched most to improve the overall visibility of a scene.

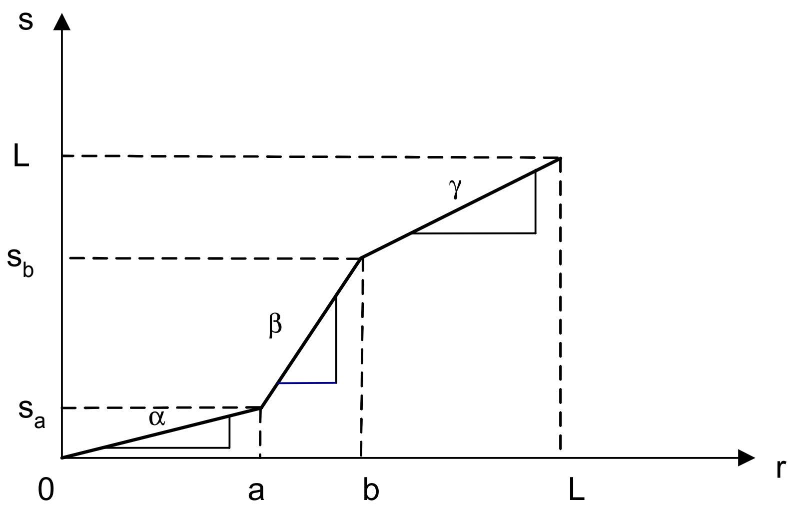

Contrast stretching transformation.
For dark region stretch : $\alpha > 1$
midregion stretch : $\beta > 1$
bright region stretch : $\gamma > 1$
(α, β, γ are gradients of the line segments)

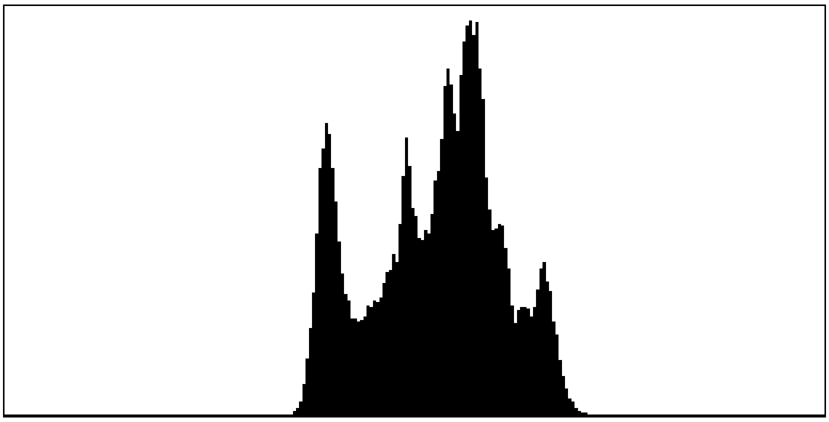

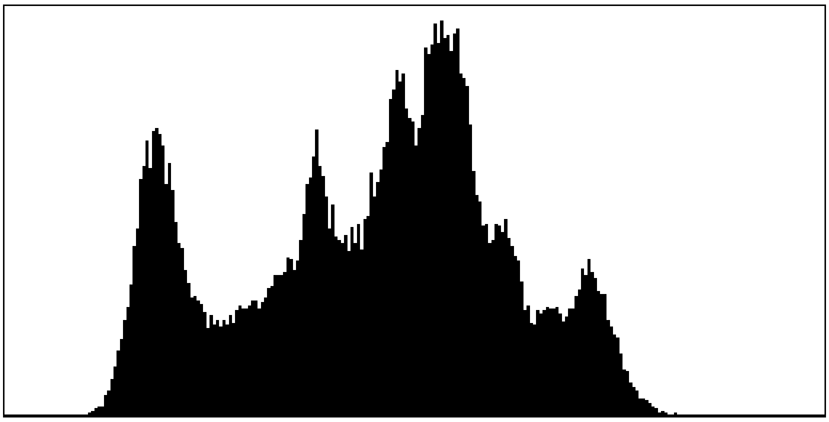

Clipping and Thresholding

In clipping, $\alpha = \gamma = 0$. This is useful for noise reduction when the input signal is known to lie in the range $[a, b]$.

Thresholding is a special case of clipping where $a = b \equiv t$ and the output becomes binary.

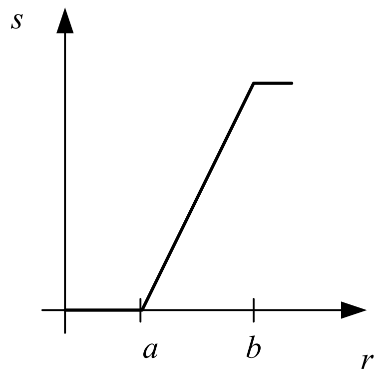

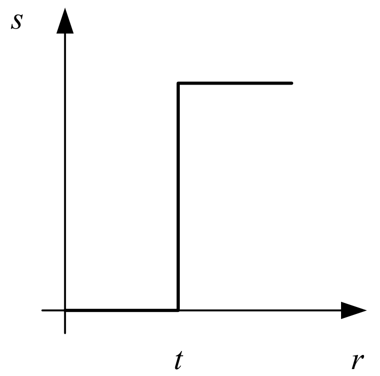

Example - Clipping

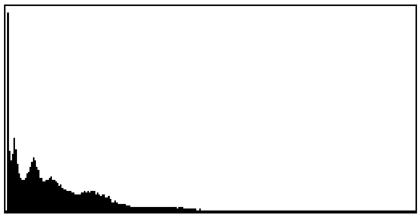

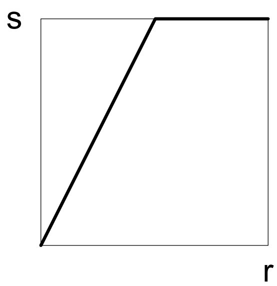

Transformation function

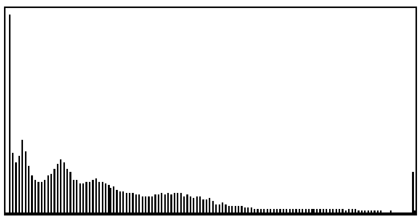

Example - Thresholding

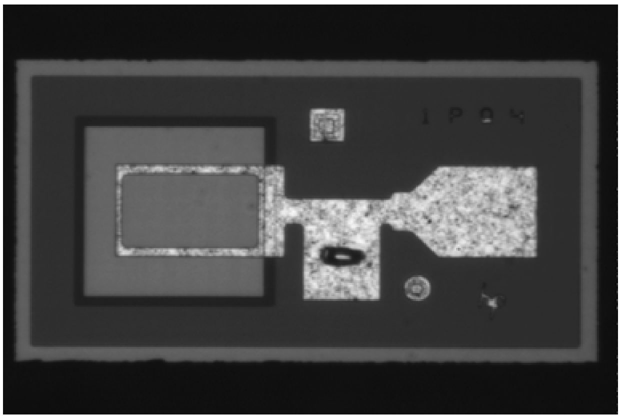

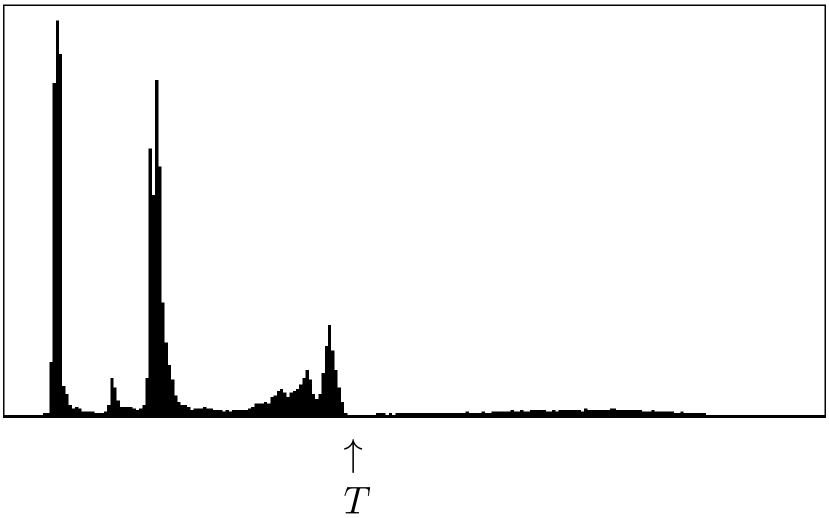

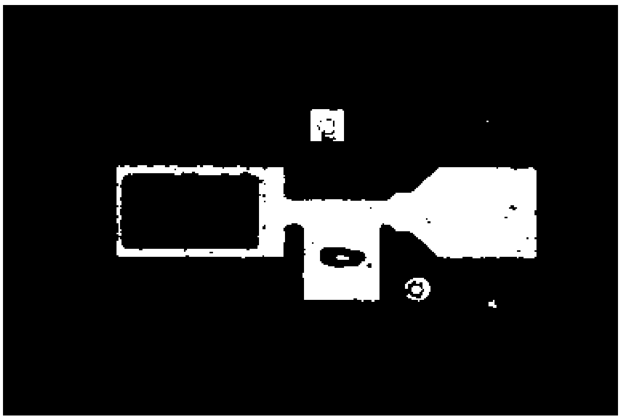

Digital Negative

A negative image can be obtained by reverse scaling of the gray levels according to the transformation

$$s = L - r \qquad (3)$$

Digital negatives are useful in the display of medical images and in producing negative prints of images.

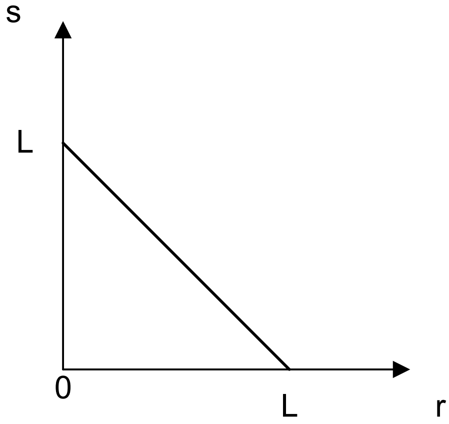

Digital negative transformation

Intensity Level Slicing

Without background:

$$s = \begin{cases} L & a \leq r \leq b \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

With background:

$$s = \begin{cases} L & a \leq r \leq b \\ r & \text{otherwise} \end{cases} \quad (5)$$

These transformations permit highlighting a specific range of gray levels in an image.

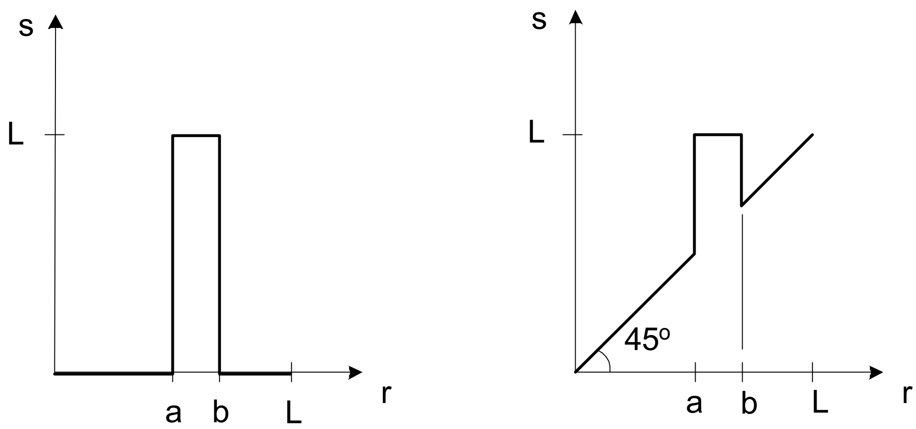

Intensity level slicing

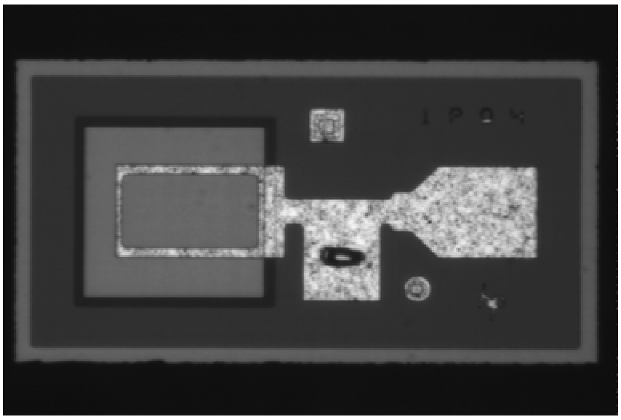

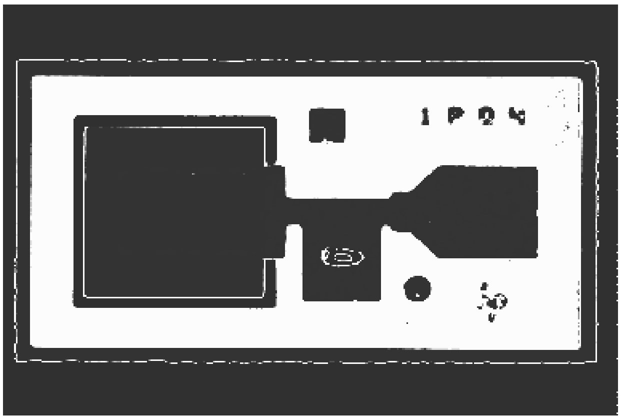

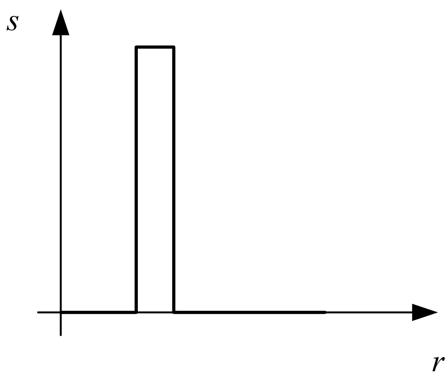
s
r

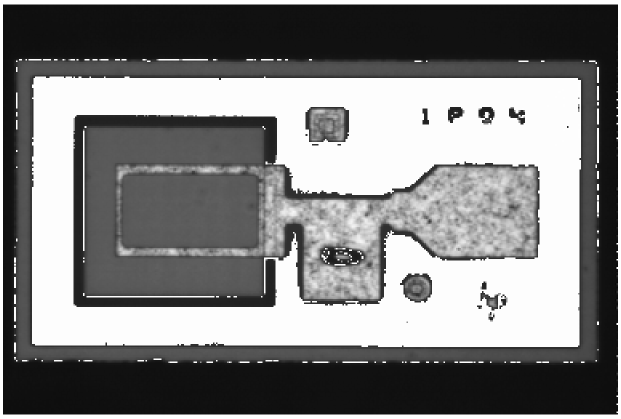

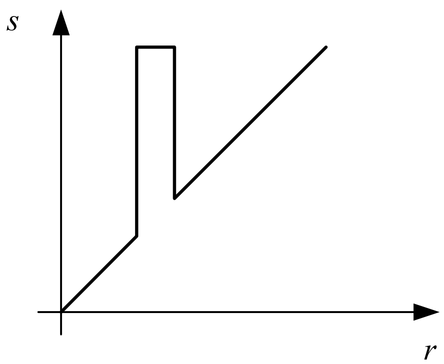
s
r

Range Compression

Sometimes the dynamic range of a processed image far exceeds the capability of the display device. An example is the display of the Fourier spectrum of an image. The dynamic range can be compressed by

$$s = c \log(1 + |r|) \tag{6}$$

where c is a scaling constant.

Consider a Fourier spectrum with values in the range $[0, 2.5 \times 10^6]$. The corresponding values of $\log(1 + |r|)$ range form 0 to 6.4. If it is desired to scale the range up to $[0, 255]$, the scaling factor is $c = 255/6.4$.

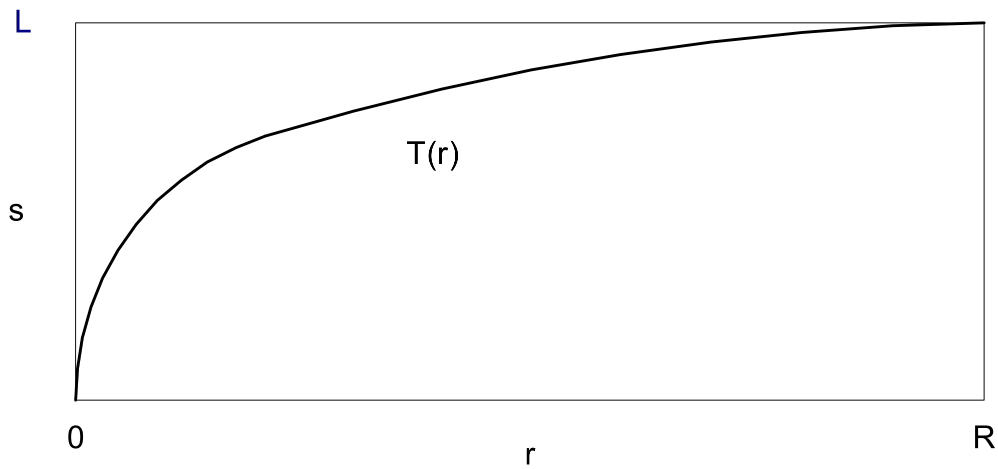

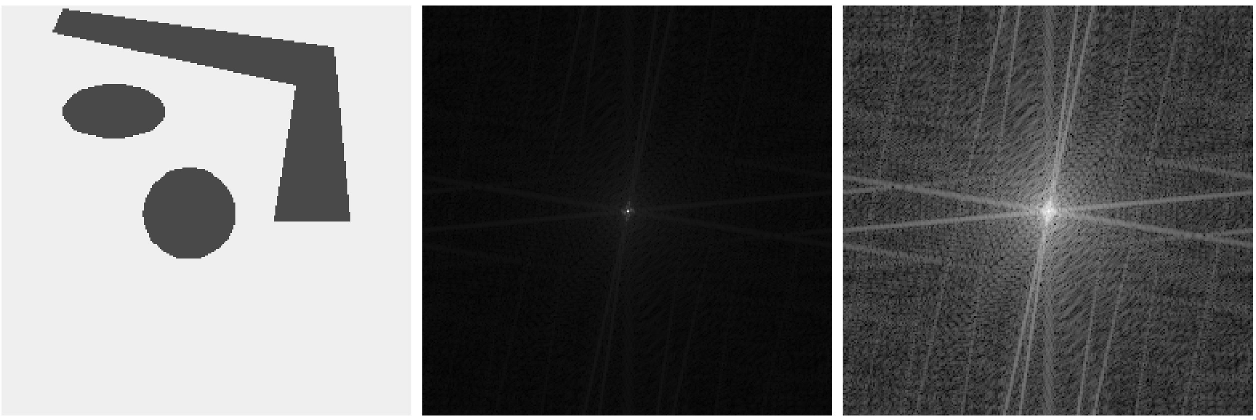

Histogram Equalization

Let r represent the intensity of pixels in an image, where r is a normalized, continuous variable lying in the range $0 \leq r \leq 1$. The transformation

$$s = T(r) \tag{7}$$

produces an intensity value s for every intensity value r in the input image. T satisfies the conditions

1. $T(r)$ is single-valued and monotonically increasing in the interval $[0, 1]$.
2. $0 \leq T(r) \leq 1$ for $0 \leq r \leq 1$.

Condition 1 preserves the order from black to white in the intensity scale; condition 2 guarantees a mapping that is consistent with the allowed 0 to 1 range of pixel values.

The inverse transformation function is:

$$r = T^{-1}(s) \tag{8}$$

where it is assumed that $T^{-1}(s)$ satisfies the two conditions above.

The intensity variables r and s are random quantities in the interval $[0, 1]$ which can be characterized by their probability density functions (PDFs) $p_r(r)$ and $p_s(s)$. Some aspects of the general appearance of an image can be deduced from its intensity PDF.

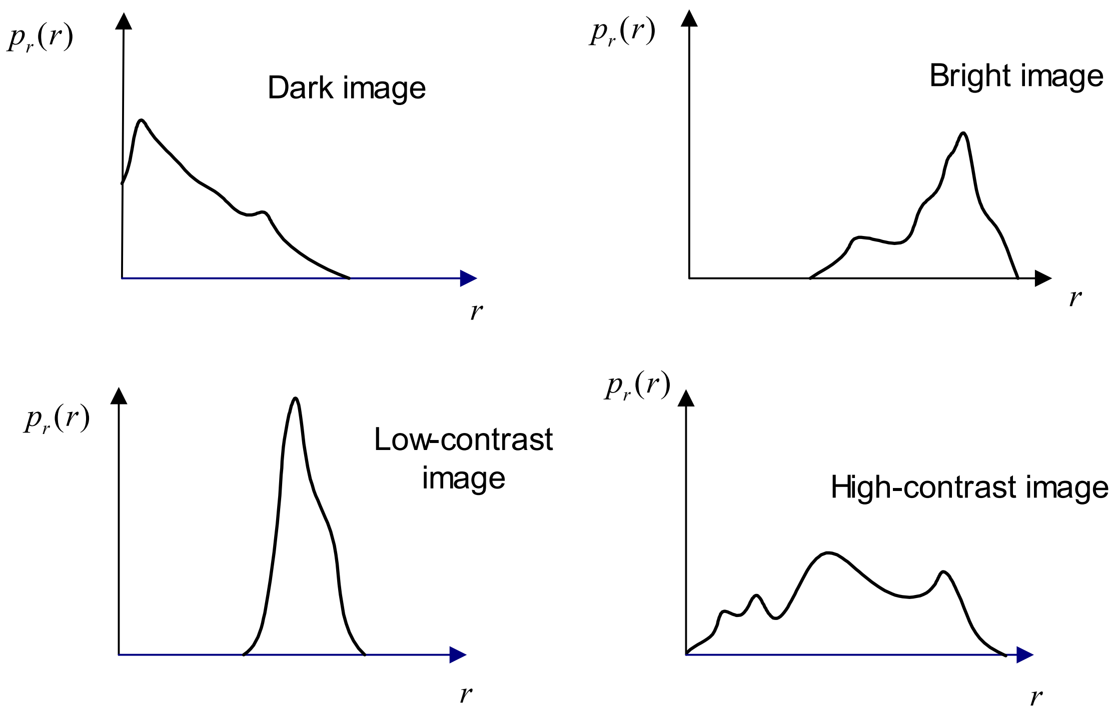

It is possible to obtain a function $T(r)$ which will transform an image such that its histogram is uniform. This process, known as histogram equalization, ensures a better utilization of all gray levels, thereby improving image contrast.

From probability theory, it follows that if $p_r(r)$ and $T(r)$ are known and $T^{-1}(s)$ satisfies condition 1, then the PDFs $p_r(r)$ and $p_s(s)$ are related by

$$p_s(s)ds = p_r(r)dr \tag{9}$$

Since we require $p_s(s) = 1$,

$$\begin{aligned} ds &= p_r(r)dr \qquad (10) \\ s &= \int p_r(r)dr \qquad (11) \\ &= \int_0^r p_r(w)dw \qquad 0 \le r \le 1 \qquad (12) \end{aligned}$$

where w is a dummy variable of integration. The rightmost side of the equation is the cumulative distribution function (CDF) of r. This transformation function satisfies the two conditions stated above.

Hence, the required transformation function is

$$s = T(r) = \int_0^r p_r(w)\, dw \qquad 0 \le r \le 1 \tag{13}$$

Note that the transformation function given in Eq. (13) yields transformed intensities that always have a flat PDF, *independent* of the shape of $p_r(r)$.

Example:

Suppose that $p_r(r)$ is given by

$$p_r(r) = \begin{cases} -2r + 2 & 0 \leq r \leq 1 \\ 0 & \text{elsewhere.} \end{cases}$$

Substitution of this expression in Eq. (13) yields

$$\begin{aligned} s &= T(r) = \int_0^r (-2w + 2)\, dw \\ &= -r^2 + 2r. \end{aligned}$$

Solving for r in terms of s yields

$$r = T^{-1}(s) = 1 \pm \sqrt{1 - s}.$$

Since r lies in the interval $[0, 1]$, only the solution

$$r = T^{-1}(s) = 1 - \sqrt{1 - s}$$

is valid.

The PDF of s is obtained by using Eq. (9):

$$\begin{aligned} p_s(s) &= \left[p_r(r) \frac{dr}{ds} \right]_{r=T^{-1}(s)} \\ &= \left[(-2r + 2) \frac{dr}{ds} \right]_{r=1-\sqrt{1-s}} \\ &= \left[(2\sqrt{1 - s} \frac{d}{ds} (1 - \sqrt{1 - s}) \right] \\ &= 1 \qquad 0 \leq s \leq 1, \end{aligned}$$

which is a uniform density in the required range.

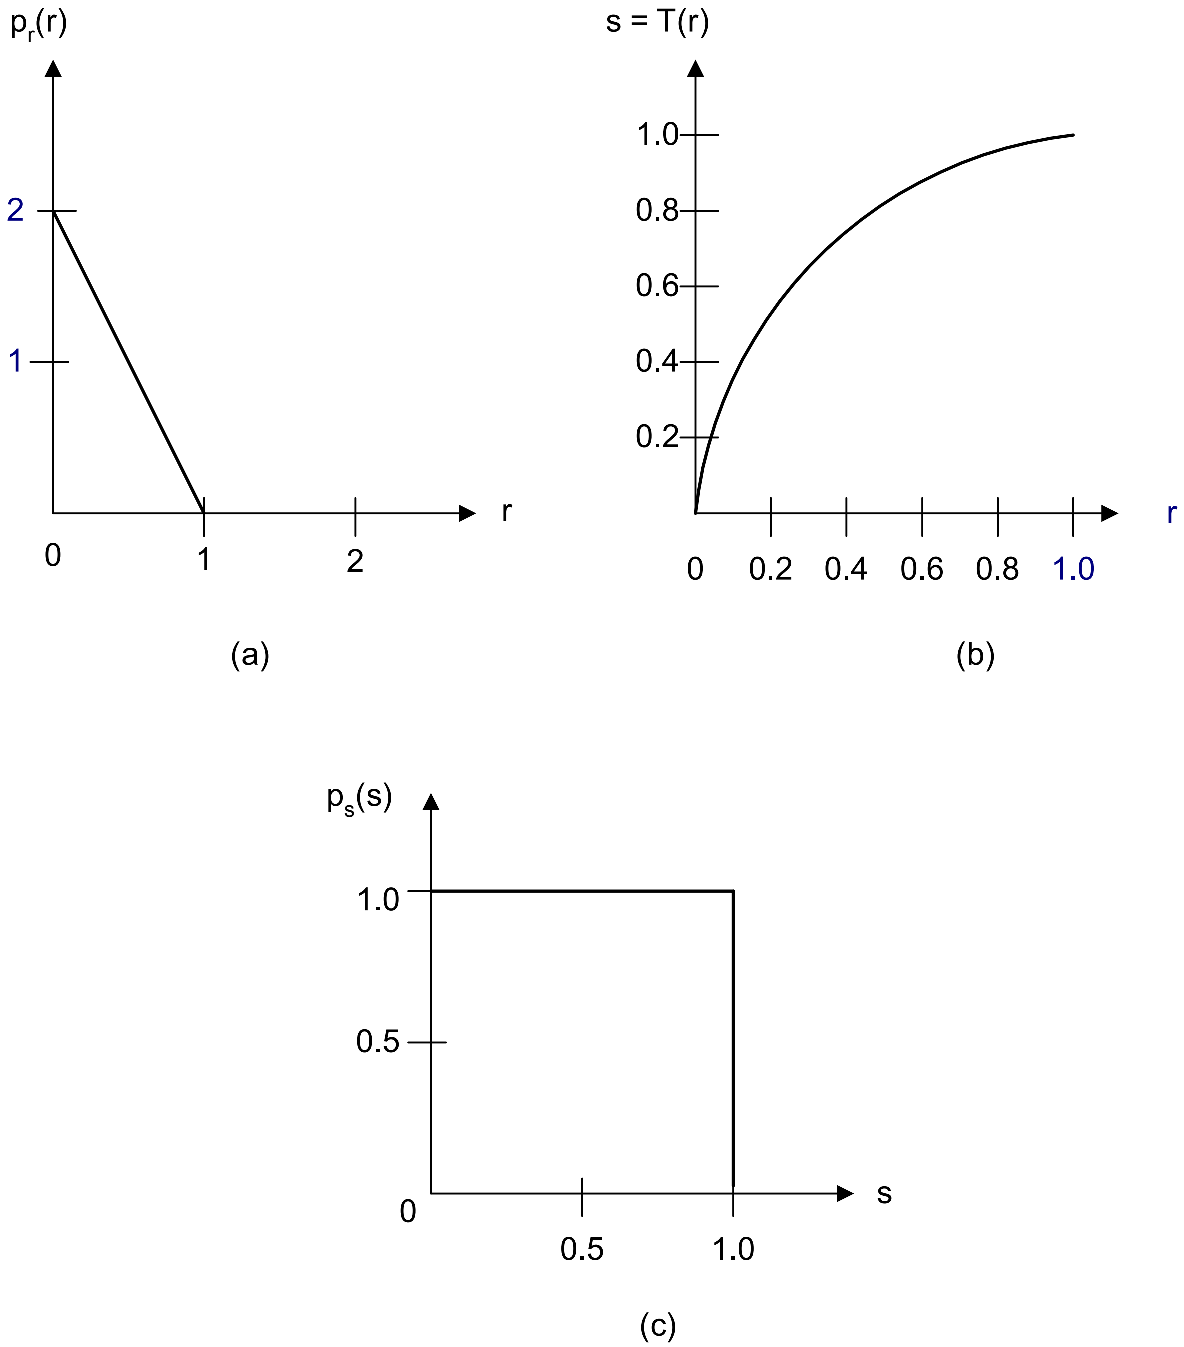

Illustration of the uniform density transformation method:
(a) orginal probability density function;
(b) transformation function;
(c) resulting uniform density.

We deal with **discrete variables** in digital image processing.

$$p_r(r_k) = \frac{n_k}{n} \tag{14}$$

where $0 \le r_k \le 1$ and $k = 0, 1, 2, \ldots, L-1$. In this equation

L	the number of discrete intensity levels
$p_r(r_k)$	an estimate of the probability of intensity r_k
n_k	the number of times this intensity appears in the image
n	the total number of pixels in the image.

A plot of $p_r(r_k)$ versus r_k is called a histogram.

The discrete form of Eq. (13) is given by

$$\begin{aligned} s_k &= T(r_k) = \sum_{j=0}^{k} \frac{n_j}{n} && (15) \\ &= \sum_{j=0}^{k} p_r(r_j) && (16) \end{aligned}$$

for $0 \le r_k \le 1$ and $k = 0, 1, 2, \ldots, L-1$.

The inverse discrete transformation is given by

$$r_k = T^{-1}(s_k) \qquad 0 \le s_k \le 1 \tag{17}$$

where both $T(r_k)$ and $T^{-1}(s_k)$ are assumed to satisfy conditions 1 and 2 above.

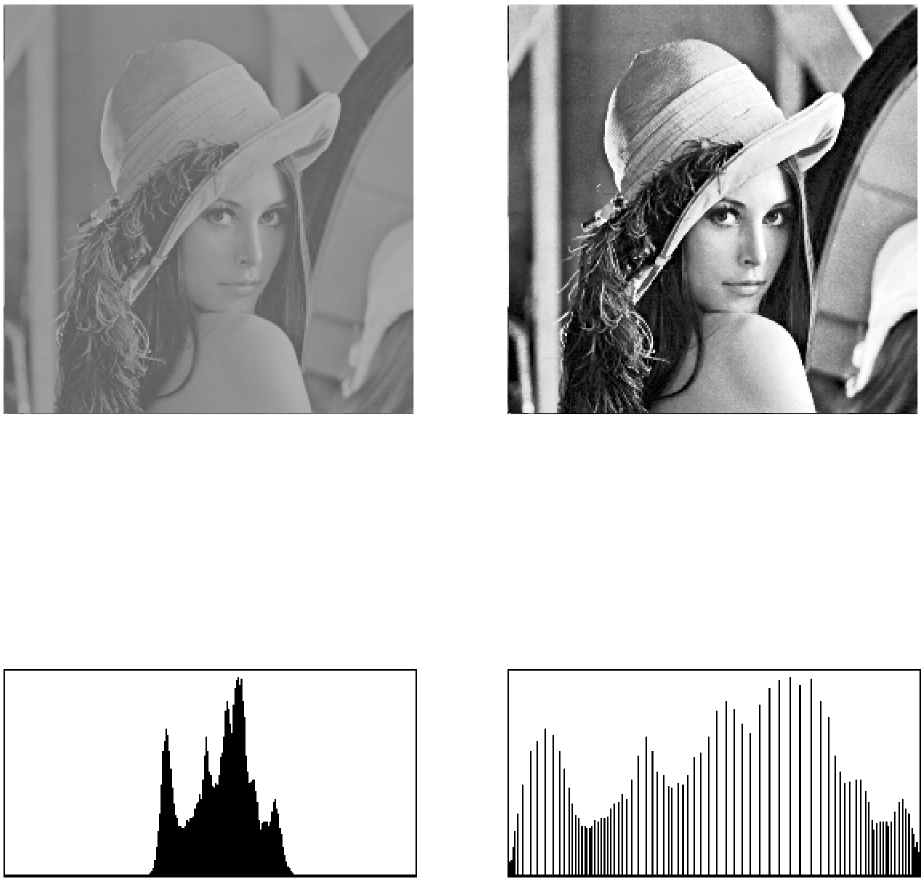

Example

Consider an 11-level input image of size 100×100.

k	r_k	n_k	$p_r(r_k) = n_k/n$
0	0	500	0.05
1	0.1	1000	0.10
2	0.2	3000	0.30
3	0.3	1500	0.15
4	0.4	500	0.05
5	0.5	1000	0.10
6	0.6	500	0.05
7	0.7	2000	0.20
8	0.8	0	0
9	0.9	0	0
10	1	0	0

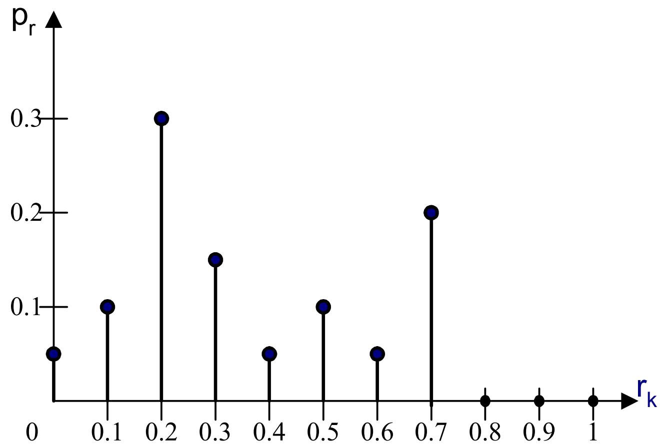

The transformation function is obtained from

$$s_k = T(r_k) = \sum_{j=0}^{k} p_r(r_j)$$

$$s_0 = T(r_0) = \sum_{j=0}^{0} p_r(r_j) = p_r(r_0) = 0.05 \to 0.1$$

$$s_1 = T(r_1) = \sum_{j=0}^{1} p_r(r_j) = p_r(r_0) + p_r(r_1) = 0.05 + 0.10 = 0.15 \to 0.2$$

$$\begin{aligned} s_2 &= 0.45 \to 0.5 \\ s_3 &= 0.60 \to 0.6 \\ s_4 &= 0.65 \to 0.7 \\ s_5 &= 0.75 \to 0.8 \\ s_6 &= 0.80 \to 0.8 \\ s_7 &= 1.00 \to 1 \\ s_8 &= 1.00 \to 1 \\ s_9 &= 1.00 \to 1 \\ s_10 &= 1.00 \to 1 \end{aligned}$$

Note: The transformed values must be assigned to the closest valid level.

r_k		s_k	$p_s(s_k)$
0	→	0.1	.05
0.1	→	0.2	.10
0.2	→	0.5	.30
0.3	→	0.6	.15
0.4	→	0.7	.05
0.5	→	0.8	.15
0.6	→	0.8	
0.7	→	1	.20
0.8	→	1	
0.9	→	1	
1	→	1	

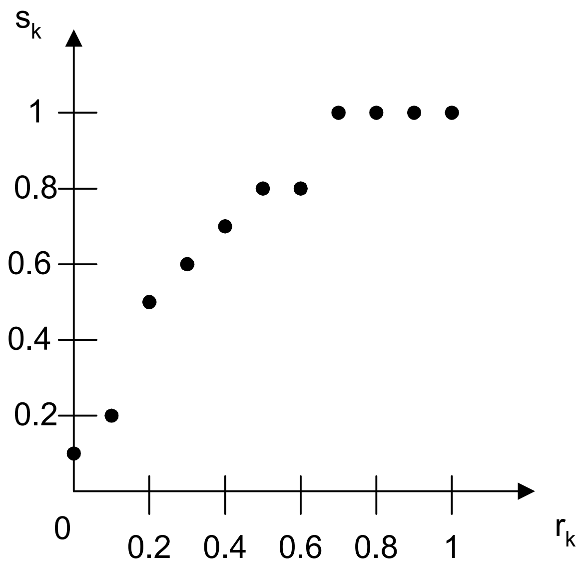

Transformation function

s_k	$p_s(s_k)$
0	0
0.1	$p_r(r_0) = 0.05$
0.2	$p_r(r_1) = 0.10$
0.3	0
0.4	0
0.5	$p_r(r_2) = 0.3$
0.6	$p_r(r_3) = 0.15$
0.7	$p_r(r_4) = 0.05$
0.8	$p_r(r_5) + p_r(r_6) = 0.15$
0.9	0
1	$p_r(r_7) + p_r(r_8) + p_r(r_9) + p_r(r_{10}) = 0.2$

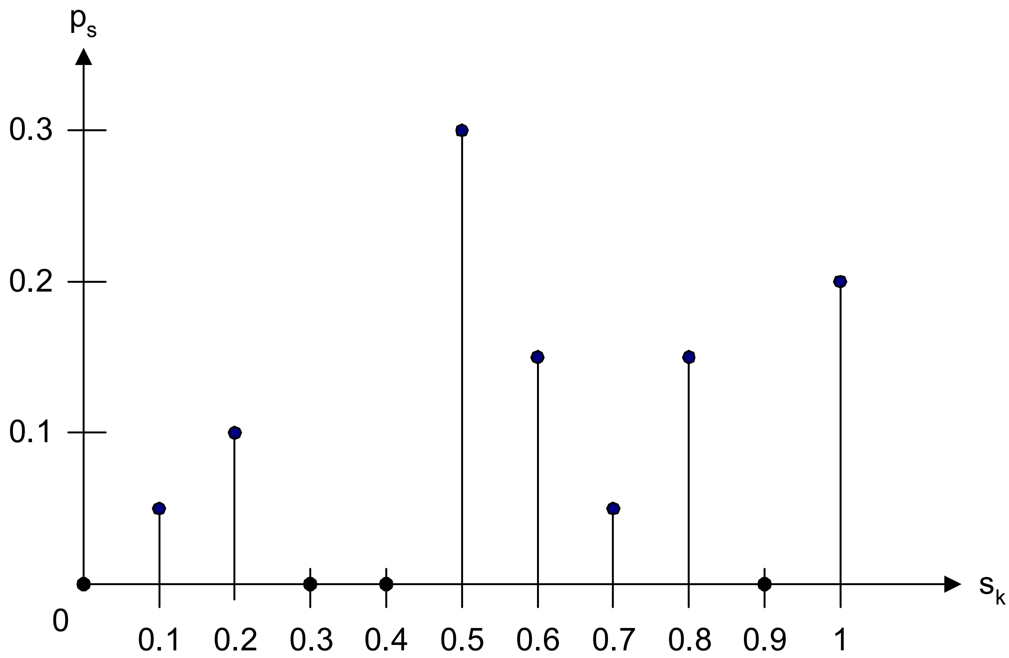

Histogram Specification

It is sometimes desirable to specify particular histograms capable of highlighting certain gray-level ranges in an image.

Let $p_r(r)$ and $p_z(z)$ be the original and desired PDFs, respectively. Suppose that a given image is first histogram equalized using Eq. (13):

$$s = T(r) = \int_0^r p_r(w)\, dw. \tag{18}$$

If the desired image were available, its levels could also be equalized by using the transformation function

$$v = G(z) = \int_0^z p_z(w)\, dw. \tag{19}$$

Note that $p_s(s)$ and $p_v(v)$ are *identical* uniform densities. Thus, instead of using v in the inverse process, we use the uniform levels s obtained from the original image, and the resulting levels, $z = G^{-1}(s)$, would have the desired PDF.

The procedure is summarized as follows:

1. Equalize the levels of the original image using Eq. (13).
2. Specify the desired density function and obtain $G(z)$ using Eq. (19).
3. Apply the inverse transformation function, $z = G^{-1}(s)$, to the levels obtained in step 1.

The two transformations required for histogram specification can be combined into a single transformation:

$$z = G^{-1}(s) = G^{-1}[T(r)] \tag{17}$$

which relates r to z.

The **discrete formulation** is as follows:

$$s_k = T(r_k) = \sum_{j=0}^{k} p_r(r_j) \tag{20}$$

$$G(z_i) = \sum_{j=0}^{i} p_z(z_j) \tag{21}$$

$$z_i = G^{-1}(s_i) \tag{22}$$

where $p_r(r_j)$ is computed from the input image and $p_z(z_j)$ is specified.

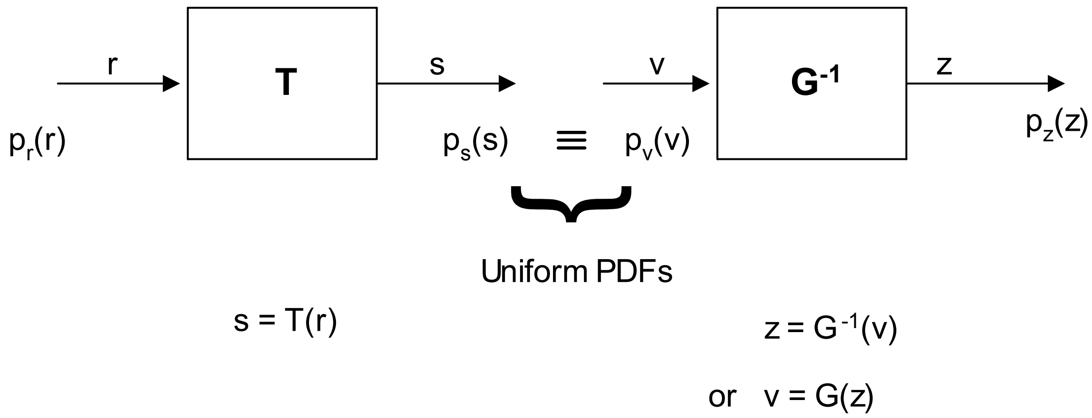

$$\begin{aligned} z &= G^{-1}(v) \\ &= G^{-1}(s) \\ &= G^{-1}[T(r)] \end{aligned}$$

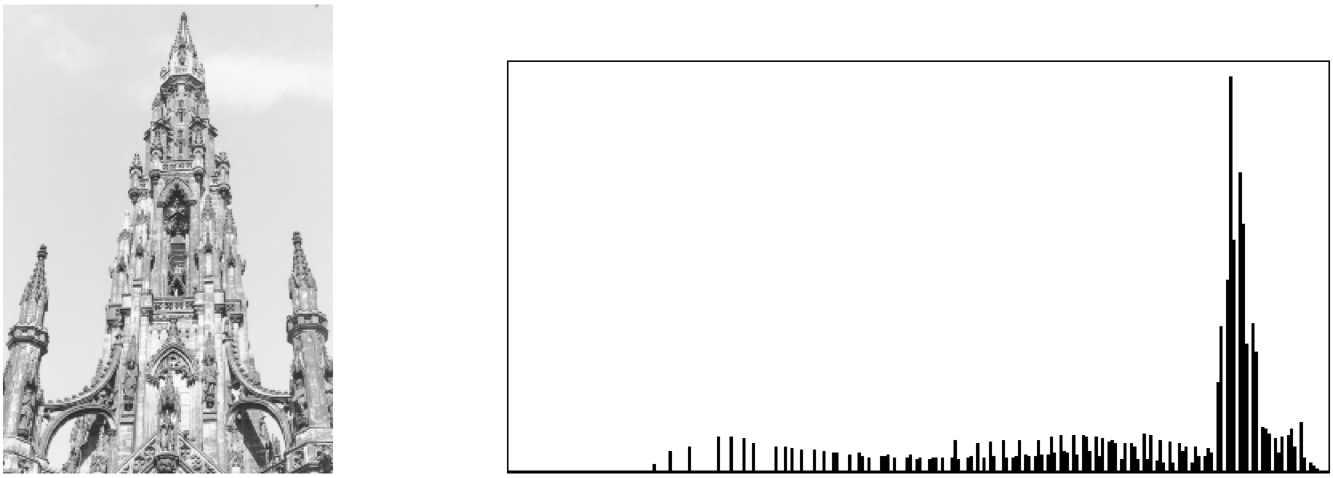

Original

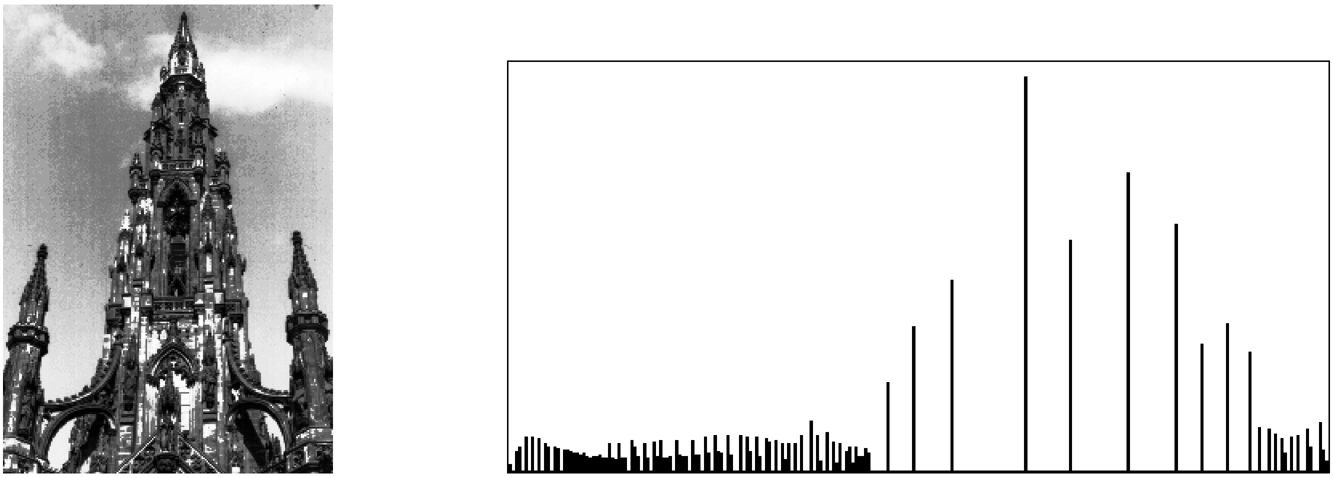

Histogram-equalised

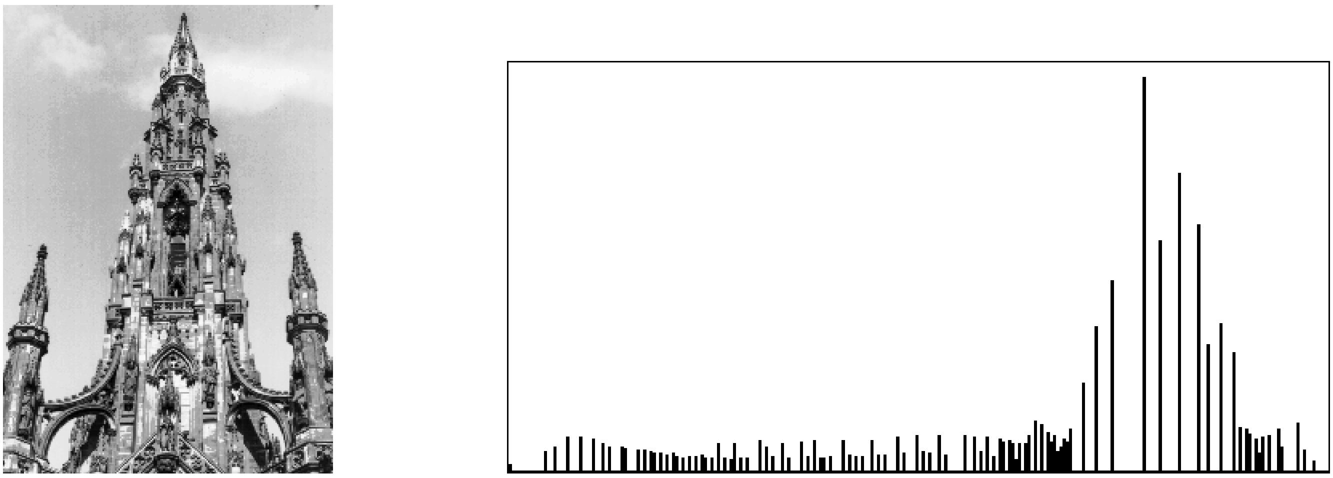

After transformation

Example

Consider a 100×100 11-level image with histogram $p_r(r_k)$. It is desired to transform the image to one with histogram $p_z(z_k)$.

r_k	$p_r(r_k)$
0	0.05
0.1	0.10
0.2	0.30
0.3	0.15
0.4	0.05
0.5	0.10
0.6	0.05
0.7	0.20
0.8	0
0.9	0
1	0

$\longrightarrow$

z_k	$p_z(z_k)$
0	0
0.1	0
0.2	0
0.3	0
0.4	0
0.5	0
0.6	0.2
0.7	0.2
0.8	0.2
0.9	0.2
1	0.2

Step 1: *Compute histogram-equalisation mappings*

(a) $r \rightarrow s$

r_k		s_k	$p_s(s_k)$
0	→	0.1	.05
0.1	→	0.2	.10
0.2	→	0.5	.30
0.3	→	0.6	.15
0.4	→	0.7	.05
0.5	→	0.8	.15
0.6	→	0.8	
0.7	→	1	.20
0.8	→	1	
0.9	→	1	
1	→	1	

(b) $z \to v$

Compute the transformation function

$$v_k = G(z_k) = \sum_{j=0}^{k} p_z(z_j)$$

$$\begin{aligned}
v_0 &= G(z_0) &= 0\\
v_1 &= G(z_1) &= 0\\
v_2 &= G(z_2) &= 0\\
v_3 &= G(z_3) &= 0\\
v_4 &= G(z_4) &= 0\\
v_5 &= G(z_5) &= 0\\
v_6 &= G(z_6) &= 0.2\\
v_7 &= G(z_7) &= 0.4\\
v_8 &= G(z_8) &= 0.6\\
v_9 &= G(z_9) &= 0.8\\
v_{10} &= G(z_{10}) &= 1
\end{aligned}$$

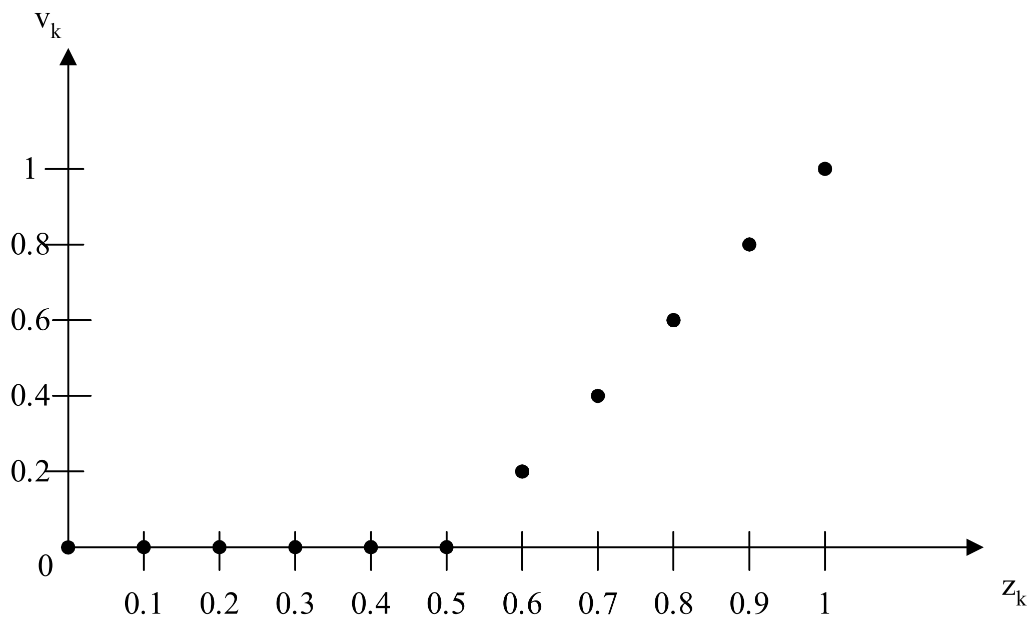

Step 2: *Compute mappings* $r \to s$, $v \to z$

r_k		s_k		z_k	$p_z(z_k)$
0	$\to$	0.1	$\to$	0.6	0.15
0.1	$\to$	0.2	$\to$	0.6	
0.2	$\to$	0.5	$\to$	0.8	0.45
0.3	$\to$	0.6	$\to$	0.8	
0.4	$\to$	0.7	$\to$	0.9	0.2
0.5, 0.6	$\to$	0.8	$\to$	0.9	
0.7, 0.8, 0.9, 1	$\to$	1	$\to$	1	0.2

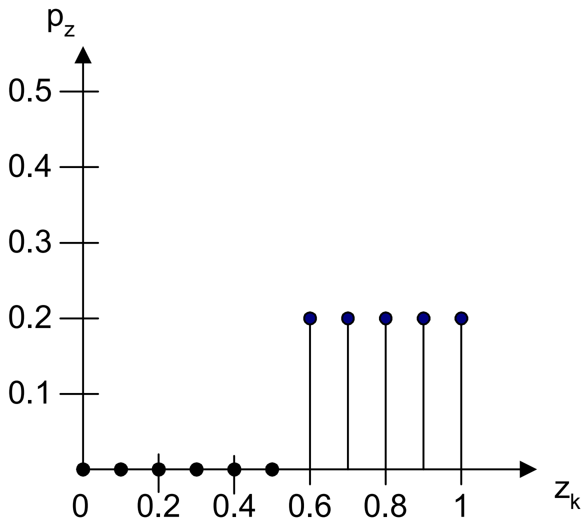

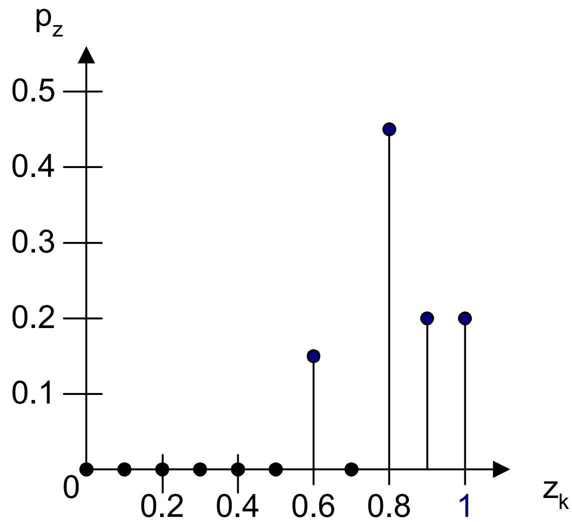

Local Enhancement

Histogram modification over a neighbourhood instead of over the entire image may be preferable in some images.

The procedure is

1. Define an $m \times n$ neighbourhood and move the centre of this neighbourhood from pixel to pixel.
2. At each location, compute the histogram of the points in the neighbourhood and obtain the transformation function.
3. Map the intensity of the centre pixel using this function.
4. Move the neighbourhood to an adjacent pixel.

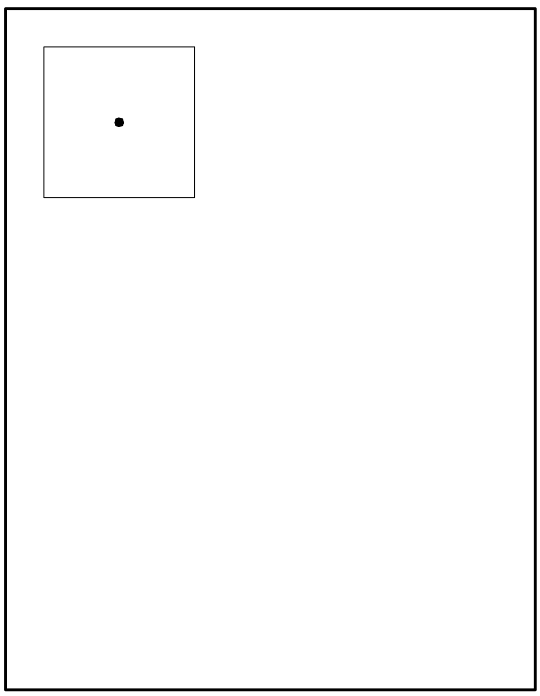

Image Sharpening

Sharpening techniques are used primarily as enhancement tools to highlight fine detail in an image or to enhance detail that has been blurred.

A sharpening operation is similar to highpass filtering in that edges, i.e., high frequency components, are boosted while low frequency components are suppressed.

A highpass filtered image may be computed as the difference between the original image and a lowpass filtered version of that image; i.e.,

$$\text{Highpass} = \text{Original} - \text{Lowpass}$$

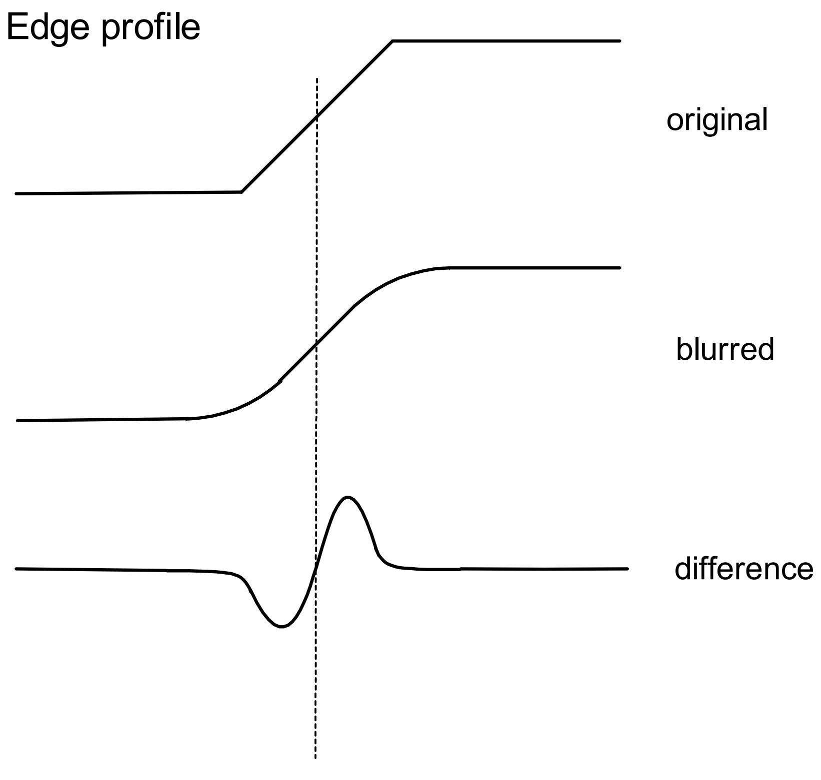

Multiplying the original image by an amplification factor A ($A \geq 1$) yields the definition of a *high-boost* or *high-frequency-emphasis* filter:

$$\begin{aligned} \text{High boost} &= (A)(\text{Original}) - \text{Lowpass} && (23) \\ &= (A-1)(\text{Original}) + \text{Original} - \text{Lowpass} && (24) \\ &= (A-1)(\text{Original}) + \text{Highpass} && (25) \end{aligned}$$

- The general process of subtracting a blurred image from an original (Eq. (23)), is called *unsharp masking.*
- $A = 1$ yields the standard highpass results.
- When $A > 1$, past of the original is added back to the highpass result, which restores partially the low-frequency components lost in the highpass filtering operation. The results is that the high-boost image looks more like the original image, with a relative degree of enhancement that depends on the value of A.

Eq. (23) may be implemented by spatial filtering with the mask

$$\begin{aligned} M &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & A & 0 \\ 0 & 0 & 0 \end{bmatrix} - \frac{1}{9} \times \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} \\ &= \frac{1}{9} \times \begin{bmatrix} -1 & -1 & -1 \\ -1 & w & -1 \\ -1 & -1 & -1 \end{bmatrix} \end{aligned}$$

where the center weight of the mask is

$$w = 9A - 1 \qquad (A \geq 1)$$

The value of A determines the nature of the filter.

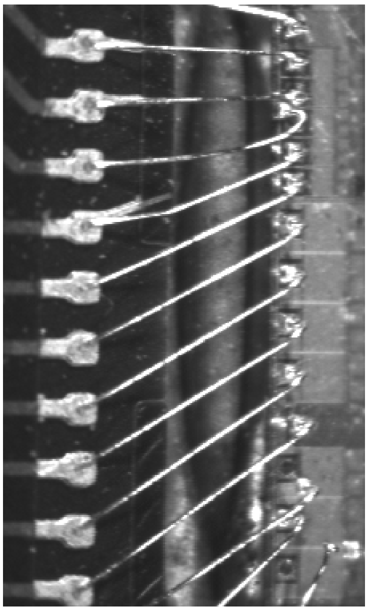

Original ("wirebond")

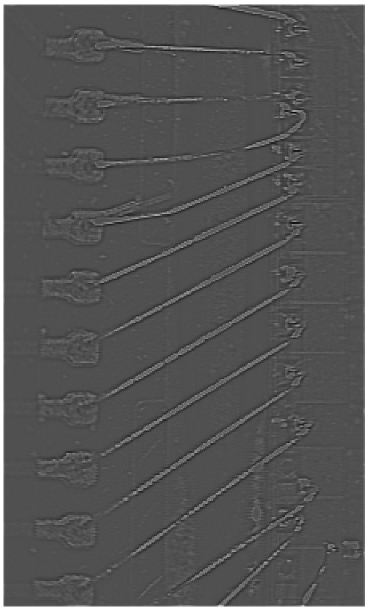

$w = 8 \quad (A = 1)$

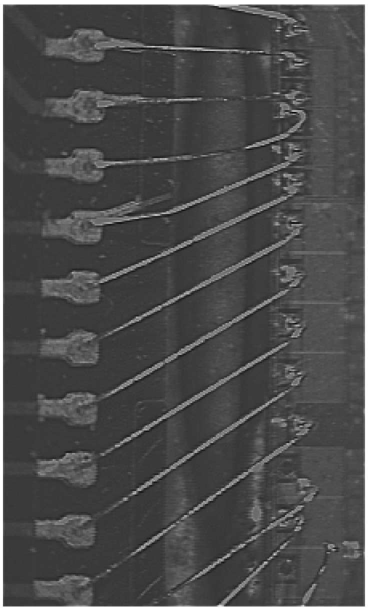

$w = 10 \quad (A = 1.22)$

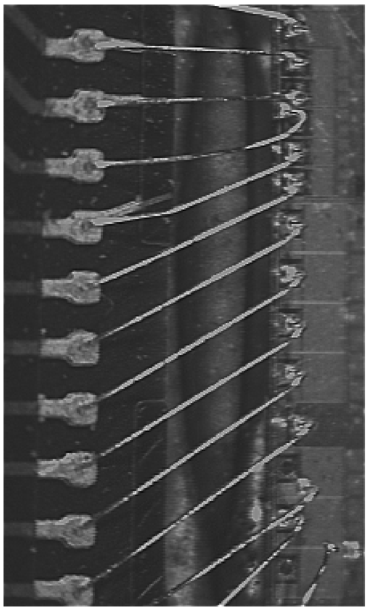

$w = 12 \quad (A = 1.44)$

(The gray levels of the processed images may be negative; hence a positive offset is applied for idsplay purposes.)

FILTERING IN THE FREQUENCY DOMAIN

Lowpass Filtering

Smoothing in achieved in the frequency domain by attenuating a specified range of high-frequency components. We have

$$G(u,v) = H(u,v)F(u,v) \tag{26}$$

The problem is to select a filter function $H(u,v)$ that yields the desired $G(u,v)$. The inverse transform of $G(u,v)$ is then the output image. We consider transfer functions that affect the real and imaginary parts of $F(u,v)$ in exactly the same manner. Such filters are referred to as zero-phase-shift filters.

Ideal Filter

A 2-D ideal low-pass filter (ILPF) has the transfer function

$$H(u,v) = \begin{cases} 1 & \text{if } D(u,v) \leq D_0 \\ 0 & \text{if } D(u,v) > D_0 \end{cases} \tag{27}$$

where D_0 is a sepcified non-negative quantity, and $D(u,v)$ is the distance from point (u,v) to the origin of the frequency plane; that is,

$$D(u,v) = (u^2 + v^2)^{1/2} \tag{28}$$

The filters considered here are radially symmetrical about the origin. Specification of radially symmetric filters centered on the $N \times N$ frequency square is based on the assumption that the origin of the Fourier transform has been centered on the square. For an ideal LPF cross section, the point of transition between $H(u,v) = 1$ and $H(u,v) = 0$ is called the cutoff frequency.

Example - Ringing

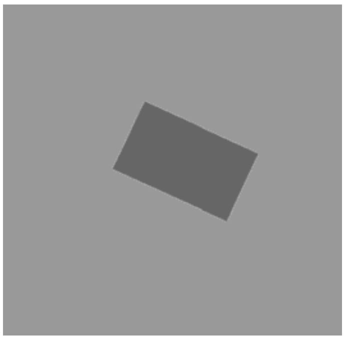

Original image

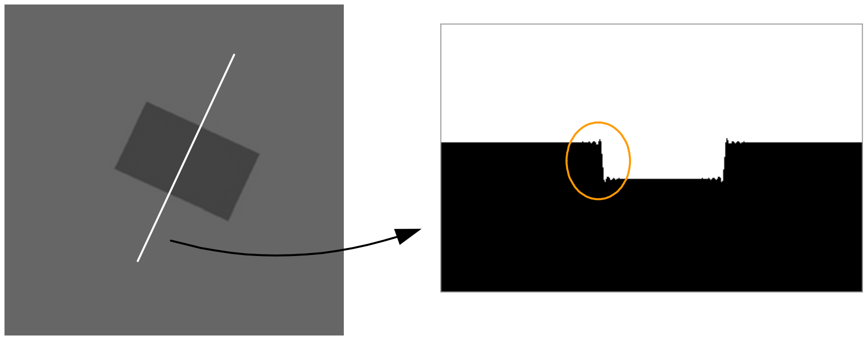

After ideal LPF

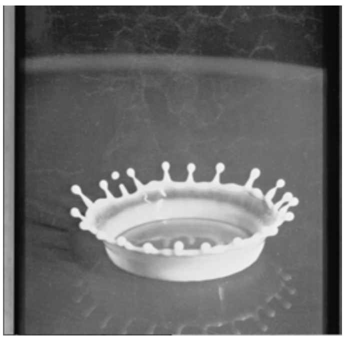

Original image

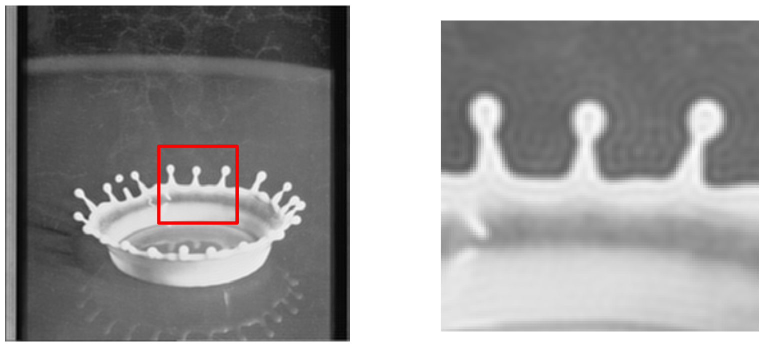

After ideal LPF

As in the 1D case, applying the ILPF results in blurring and ringing. This can be explained by first noting that the output image is

$$g(x, y) = h(x, y) \star f(x, y) \tag{29}$$

where $h(x, y)$ is the inverse Fourier transform of the filter transfer function $H(u, v)$.

For an ILPF, $h(x, y)$ has the general form of a 2D sinc function, i.e.,

$$h(x, y) \sim \text{sinc}(ax)\text{sinc}(by) \tag{30}$$

Thus convolving the input image with $h(x, y)$ will result in ringing which is more noticeable around prominent bright points or lines.

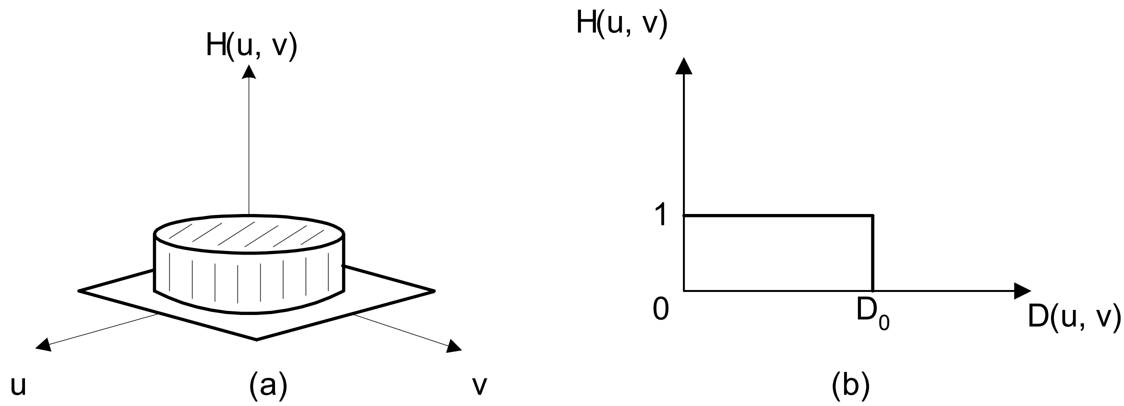

(a) Perspective plot of an ideal lowpass filter transfer function;
(b) filter cross section.

Butterworth Filter

The transfer function of the Butterworth LPF (BPLF) of order n and with cutoff frequency locus a distance D_0 from the origin is defined by the relation

$$H(u,v) = \frac{1}{1 + [D(u,v)/D_0]^{2n}} \tag{31}$$

$H(u,v)$ has value 0.5 when $D(u,v) = D_0$.

The BLPF does not have a sharp discontinuity that establishes a clear cutoff bewteen passed and filtered frequencies. The filtered image will not exhibit any ringing.

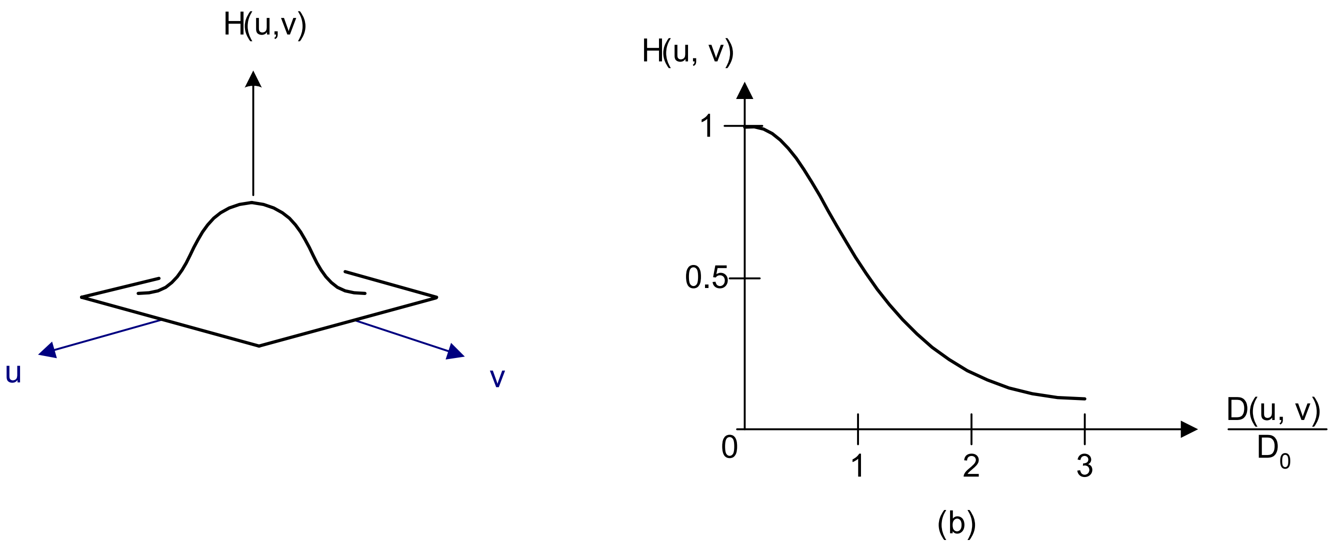

(a) A Butterworth lowpass filter; (b) radial cross section for n = 1.

Example - Low-pass filtering

Original

Noisy

Filtered using ideal LPF

Filtered using Butterworth LPF

Highpass Filtering

The ideal high-pass filter (IHPF) is defined by the transfer function

$$H(u,v) = \begin{cases} 0 & \text{if } D(u,v) \leq D_0 \\ 1 & \text{if } D(u,v) > D_0 \end{cases} \tag{32}$$

This filter is the opposite of the ILPF because it completely attenuates all frequencies inside a circle of radius D_0 while passing, without attenuation, all frequencies outside the circle.

The transfer function of the Butterworth high-pass filter (BHPF) of order n and cutoff frequency a distance D_0 from the origin is

$$H(u,v) = \frac{1}{1 + [D_0/D(u,v)]^{2n}} \tag{33}$$

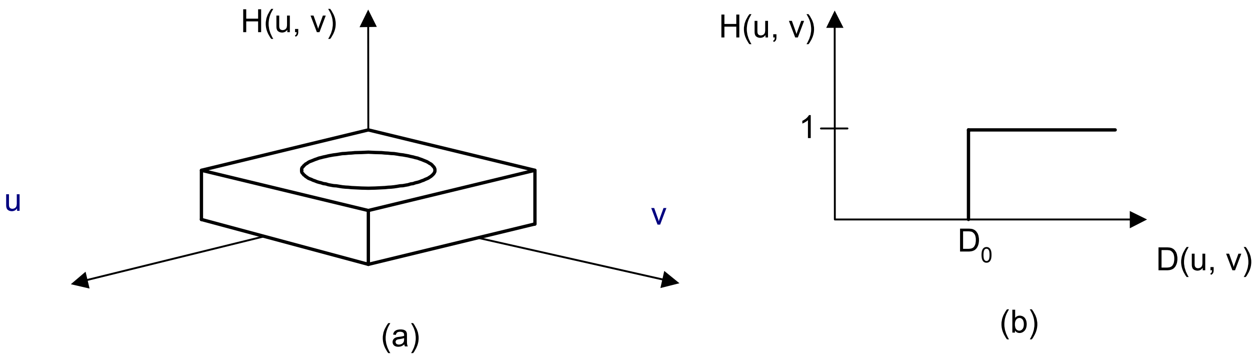

Perspective plot and radial cross section of ideal highpass filter.

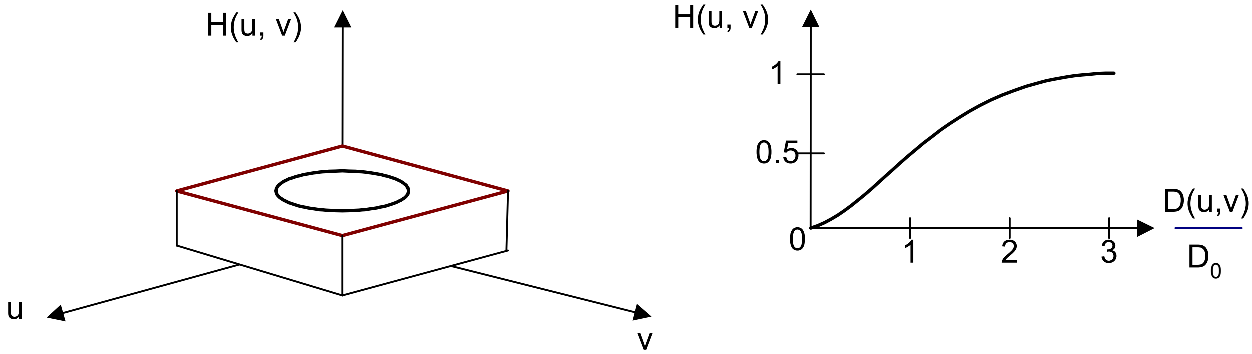

Perspective plot and radial cross section of Butterworth highpass filter for n = 1

Example - High-pass filtering

(a) Original

(b) Filtered using ideal HPF

(c) After thresholding (b)

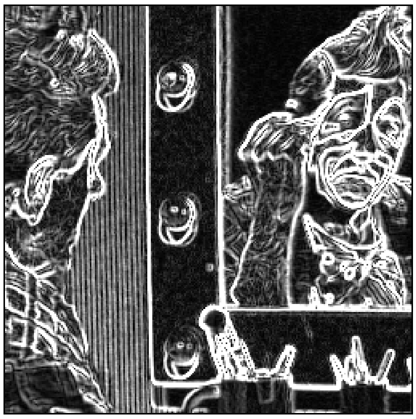

(d) After applying Sobel to (c)

Other Uses of Filtering

Frequency filters are quite useful when processing parts of an image which can be associated with certain frequencies. For example, in the figure, each part of the house is made of stripes of a different frequency and orientation.

From the Fourier transform, we can see the main peaks in the image corresponding to the periodic patterns in the spatial domain image which now can be accessed separately. For example, we can smooth the vertical stripes (i.e. those components which make up the wall in the spatial image) by multiplying the Fourier image with an appropriate filter (as shown).

The effect is that all frequencies within the black rectangle are set to zero, the others remain unchanged. Applying the inverse Fourier Transform and normalizing the resulting image gives the result shown.

Although the image shows some regular patterns in the formerly constant background, the vertical stripes are almost totally removed whereas the other patterns remained mostly unchanged.

Example

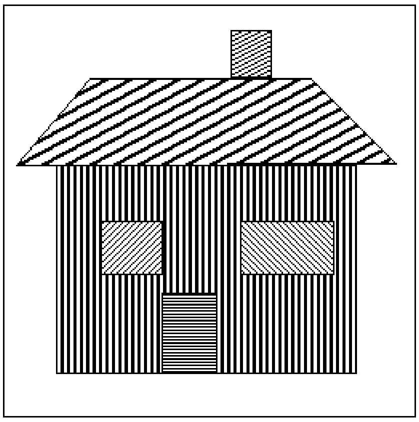

(a) Original

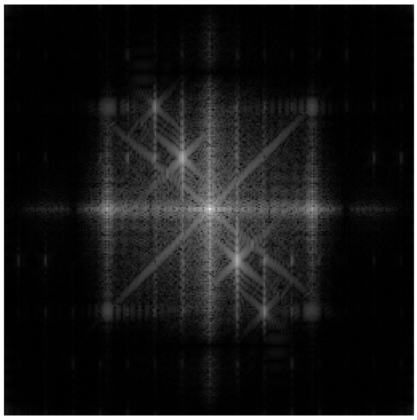

(b) Fourier transform magnitude

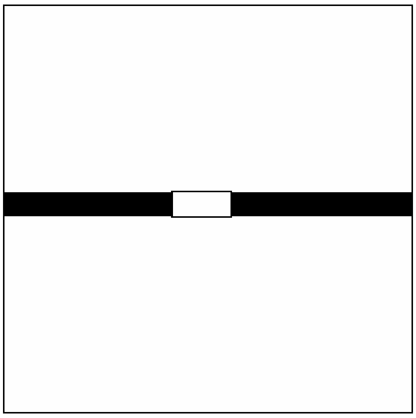

(c) Filter function

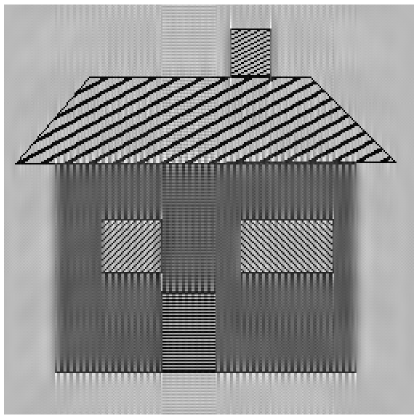

(d) After filtering (vertical lines removed)