

1. Explain about gray level interpolation.

The distortion correction equations yield non integer values for x' and y' . Because the distorted image g is digital, its pixel values are defined only at integer coordinates. Thus using non integer values for x' and y' causes a mapping into locations of g for which no gray levels are defined. Inferring what the gray-level values at those locations should be, based only on the pixel values at integer coordinate locations, then becomes necessary. The technique used to accomplish this is called gray-level interpolation.

The simplest scheme for gray-level interpolation is based on a nearest neighbor approach. This method, also called zero-order interpolation, is illustrated in Fig. 6.1. This figure shows

(A) The mapping of integer (x, y) coordinates into fractional coordinates (x', y') by means of following equations

$$x' = c_1x + c_2y + c_3xy + c_4$$

and

$$y' = c_5x + c_6y + c_7xy + c_8$$

(B) The selection of the closest integer coordinate neighbor to (x', y') ;

and

(C) The assignment of the gray level of this nearest neighbor to the pixel located at (x, y) .




Fig. 6.1 Gray-level interpolation based on the nearest neighbor concept.

Although nearest neighbor interpolation is simple to implement, this method often has the drawback of producing undesirable artifacts, such as distortion of straight edges in images of high resolution. Smoother results can be obtained by using more sophisticated techniques, such as cubic convolution interpolation, which fits a surface of the $\sin(z)/z$ type through a much larger number of neighbors (say, 16) in order to obtain a smooth estimate of the gray level at any

desired point. Typical areas in which smoother approximations generally are required include 3-D graphics and medical imaging. The price paid for smoother approximations is additional computational burden. For general-purpose image processing a bilinear interpolation approach that uses the gray levels of the four nearest neighbors usually is adequate. This approach is straightforward. Because the gray level of each of the four integral nearest neighbors of a non integral pair of coordinates (x', y') is known, the gray-level value at these coordinates, denoted $v(x', y')$, can be interpolated from the values of its neighbors by using the relationship

$$v(x', y') = ax' + by' + c x' y' + d$$

where the four coefficients are easily determined from the four equations in four unknowns that can be written using the four known neighbors of (x', y') . When these coefficients have been determined, $v(x', y')$ is computed and this value is assigned to the location in $f\{x, y\}$ that yielded the spatial mapping into location (x', y') . It is easy to visualize this procedure with the aid of Fig. 6.1. The exception is that, instead of using the gray-level value of the nearest neighbor to (x', y') , we actually interpolate a value at location (x', y') and use this value for the gray-level assignment at (x, y) .

2. Explain about Wiener filter used for image restoration.

The inverse filtering approach makes no explicit provision for handling noise. This approach incorporates both the degradation function and statistical characteristics of noise into the restoration process. The method is founded on considering images and noise as random processes, and the objective is to find an estimate f of the uncorrupted image f such that the mean square error between them is minimized. This error measure is given by

$$e^2 = E \{(f - \hat{f})^2\}$$

where $E\{\cdot\}$ is the expected value of the argument. It is assumed that the noise and the image are uncorrelated; that one or the other has zero mean; and that the gray levels in the estimate are a linear function of the levels in the degraded image. Based on these conditions, the minimum of the error function is given in the frequency domain by the expression

$$\begin{aligned}
 \hat{F}(u, v) &= \left[\frac{H^*(u, v)S_f(u, v)}{S_f(u, v)|H(u, v)|^2 + S_\eta(u, v)} \right] G(u, v) \\
 &= \left[\frac{H^*(u, v)}{|H(u, v)|^2 + S_\eta(u, v)/S_f(u, v)} \right] G(u, v) \\
 &= \left[\frac{1}{H(u, v)} \frac{|H(u, v)|^2}{|H(u, v)|^2 + S_\eta(u, v)/S_f(u, v)} \right] G(u, v)
 \end{aligned}$$

where we used the fact that the product of a complex quantity with its conjugate is equal to the magnitude of the complex quantity squared. This result is known as the Wiener filter, after N. Wiener [1942], who first proposed the concept in the year shown. The filter, which consists of the terms inside the brackets, also is commonly referred to as the minimum mean square error filter or the least square error filter. The Wiener filter does not have the same problem as the inverse filter with zeros in the degradation function, unless both $H(u, v)$ and $S_\eta(u, v)$ are zero for the same value(s) of u and v .

The terms in above equation are as follows:

$H(u, v)$ = degradation function

$H^*(u, v)$ = complex conjugate of $H(u, v)$

$|H(u, v)|^2 = H^*(u, v) * H(u, v)$

$S_\eta(u, v) = |N(u, v)|^2 =$ power spectrum of the noise

$S_f(u, v) = |F(u, v)|^2 =$ power spectrum of the undegraded image.

As before, $H(u, v)$ is the transform of the degradation function and $G(u, v)$ is the transform of the degraded image. The restored image in the spatial domain is given by the inverse Fourier transform of the frequency-domain estimate $F(u, v)$. Note that if the noise is zero, then the noise power spectrum vanishes and the Wiener filter reduces to the inverse filter.

When we are dealing with spectrally white noise, the spectrum $|N(u, v)|^2$ is a constant, which simplifies things considerably. However, the power spectrum of the undegraded image seldom is known. An approach used frequently when these quantities are not known or cannot be estimated is to approximate the equation as

$$\hat{F}(u, v) = \left[\frac{1}{H(u, v)} \frac{|H(u, v)|^2}{|H(u, v)|^2 + K} \right] G(u, v)$$

where K is a specified constant.

3. Explain a Model of the Image Degradation/Restoration Process.

The Fig. 6.3 shows, the degradation process is modeled as a degradation function that, together with an additive noise term, operates on an input image $f(x, y)$ to produce a degraded image $g(x, y)$. Given $g(x, y)$, some knowledge about the degradation function H , and some knowledge about the additive noise term $\eta(x, y)$, the objective of restoration is to obtain an estimate $\hat{f}(x, y)$ of the original image. the estimate should be as close as possible to the original input image and, in general, the more we know about H and η , the closer $\hat{f}(x, y)$ will be to $f(x, y)$.

The degraded image is given in the spatial domain by

$$g(x, y) = h(x, y) * f(x, y) + \eta(x, y)$$

where $h(x, y)$ is the spatial representation of the degradation function and, the symbol $*$ indicates convolution. Convolution in the spatial domain is equal to multiplication in the frequency domain, hence

$$G(u, v) = H(u, v) F(u, v) + N(u, v)$$

where the terms in capital letters are the Fourier transforms of the corresponding terms in above equation.




Fig. 6.3 model of the image degradation/restoration process.

4. Explain about the restoration filters used when the image degradation is due to noise only.

If the degradation present in an image is only due to noise, then,

$$g(x, y) = f(x, y) + \eta(x, y)$$

$$G(u, v) = F(u, v) + N(u, v)$$

The restoration filters used in this case are,

1. Mean filters
2. Order static filters and
3. Adaptive filters

Also read 5, 6, 7 answers.

5. Explain Mean filters.

There are four types of mean filters. They are

(i) Arithmetic mean filter

This is the simplest of the mean filters. Let S_{xy} represent the set of coordinates in a rectangular subimage window of size $m \times n$, centered at point (x, y) . The arithmetic mean filtering process computes the average value of the corrupted image $g(x, y)$ in the area defined by S_{xy} . The value of the restored image f at any point (x, y) is simply the arithmetic mean computed using the pixels in the region defined by S_{xy} . In other words

$$\hat{f}(x, y) = \frac{1}{mn} \sum_{(s,t) \in S_{xy}} g(s, t).$$

This operation can be implemented using a convolution mask in which all coefficients have value $1/mn$

(ii) Geometric mean filter

An image restored using a geometric mean filter is given by the expression

$$\hat{f}(x, y) = \left[\prod_{(s,t) \in S_{xy}} g(s, t) \right]^{\frac{1}{mn}}$$

Here, each restored pixel is given by the product of the pixels in the subimage window, raised to the power $1/mn$. A geometric mean filter achieves smoothing comparable to the arithmetic mean filter, but it tends to lose less image detail in the process.

(iii) Harmonic mean filter

The harmonic mean filtering operation is given by the expression

$$\hat{f}(x, y) = \frac{mn}{\sum_{(s,t) \in S_{xy}} \frac{1}{g(s, t)}}$$

The harmonic mean filter works well for salt noise, but fails for pepper noise. It does well also with other types of noise like Gaussian noise.

(iv) Contra harmonic mean filter

The contra harmonic mean filtering operation yields a restored image based on the expression

$$\hat{f}(x, y) = \frac{\sum_{(s,t) \in S_{xy}} g(s, t)^{Q+1}}{\sum_{(s,t) \in S_{xy}} g(s, t)^Q}$$

where Q is called the order of the filter. This filter is well suited for reducing or virtually eliminating the effects of salt-and-pepper noise. For positive values of Q , the filter eliminates pepper noise. For negative values of Q it eliminates salt noise. It cannot do both simultaneously. Note that the contra harmonic filter reduces to the arithmetic mean filter if $Q = 0$, and to the harmonic mean filter if $Q = -1$.

6. Explain the Order-Statistic Filters.

There are four types of Order-Statistic filters. They are

(i) Median filter

The best-known order-statistics filter is the median filter, which, as its name implies, replaces the value of a pixel by the median of the gray levels in the neighborhood of that pixel:

$$\hat{f}(x, y) = \text{median}_{(s,t) \in S_{xy}} \{g(s, t)\}.$$

The original value of the pixel is included in the computation of the median. Median filters are quite popular because, for certain types of random noise, they provide excellent noise-reduction capabilities, with considerably less blurring than linear smoothing filters of similar size. Median filters are particularly effective in the presence of both bipolar and unipolar impulse noise.

(ii) Max and min filters

Although the median filter is by far the order-statistics filter most used in image processing, it is by no means the only one. The median represents the 50th percentile of a ranked set of numbers, but the reader will recall from basic statistics that ranking lends itself to many other possibilities. For example, using the 100th percentile results in the so-called max filter, given by

$$\hat{f}(x, y) = \max_{(s,t) \in S_{xy}} \{g(s, t)\}.$$

This filter is useful for finding the brightest points in an image. Also, because pepper noise has very low values, it is reduced by this filter as a result of the max selection process in the subimage area S_{xy} .

The 0th percentile filter is the min filter.

$$\hat{f}(x, y) = \min_{(s,t) \in S_{xy}} \{g(s, t)\}.$$

This filter is useful for finding the darkest points in an image. Also, it reduces salt noise as a result of the min operation.

(iii) Midpoint filter

The midpoint filter simply computes the midpoint between the maximum and minimum values in the area encompassed by the filter:

$$\hat{f}(x, y) = \frac{1}{2} \left[\max_{(s,t) \in S_{xy}} \{g(s, t)\} + \min_{(s,t) \in S_{xy}} \{g(s, t)\} \right].$$

Note that this filter combines order statistics and averaging. This filter works best for randomly distributed noise, like Gaussian or uniform noise.

(iv) Alpha - trimmed mean filter

It is a filter formed by deleting the $d/2$ lowest and the $d/2$ highest gray-level values of $g(s, t)$ in the neighborhood S_{xy} . Let $g_r(s, t)$ represent the remaining $mn - d$ pixels. A filter formed by averaging these remaining pixels is called an alpha-trimmed mean filter:

$$\hat{f}(x, y) = \frac{1}{mn - d} \sum_{(s,t) \in S_{xy}} g_r(s, t)$$

where the value of d can range from 0 to $mn - 1$. When $d = 0$, the alpha- trimmed filter reduces to the arithmetic mean filter. If $d = (mn - 1)/2$, the filter becomes a median filter. For other values of d , the alpha-trimmed filter is useful in situations involving multiple types of noise, such as a combination of salt-and-pepper and Gaussian noise.

7. Explain the Adaptive Filters.

Adaptive filters are filters whose behavior changes based on statistical characteristics of the image inside the filter region defined by the $m \times n$ rectangular window S_{xy} .

Adaptive, local noise reduction filter:

The simplest statistical measures of a random variable are its mean and variance. These are reasonable parameters on which to base an adaptive filter because they are quantities closely related to the appearance of an image. The mean gives a measure of average gray level in the region over which the mean is computed, and the variance gives a measure of average contrast in that region.

This filter is to operate on a local region, S_{xy} . The response of the filter at any point (x, y) on which the region is centered is to be based on four quantities: (a) $g(x, y)$, the value of the noisy image at (x, y) ; (b) σ^2 , the variance of the noise corrupting $g(x, y)$ to form $g(x, y)$; (c) μ , the local mean of the pixels in S_{xy} ; and (d) σ^2_L , the local variance of the pixels in S_{xy} .

The behavior of the filter to be as follows:

1. If σ_η^2 is zero, the filter should return simply the value of $g(x, y)$. This is the trivial, zero-noise case in which $g(x, y)$ is equal to $f(x, y)$.
2. If the local variance is high relative to σ_η^2 the filter should return a value close to $g(x, y)$. A high local variance typically is associated with edges, and these should be preserved.
3. If the two variances are equal, we want the filter to return the arithmetic mean value of the pixels in S_{xy} . This condition occurs when the local area has the same properties as the overall image, and local noise is to be reduced simply by averaging.

Adaptive local noise filter is given by,

$$\hat{f}(x, y) = g(x, y) - \frac{\sigma_\eta^2}{\sigma_L^2} [g(x, y) - m_L].$$

The only quantity that needs to be known or estimated is the variance of the overall noise, σ_η^2 . The other parameters are computed from the pixels in S_{xy} at each location (x, y) on which the filter window is centered.

Adaptive median filter:

The median filter performs well as long as the spatial density of the impulse noise is not large (as a rule of thumb, P_a and P_b less than 0.2). The adaptive median filtering can handle impulse noise with probabilities even larger than these. An additional benefit of the adaptive median filter is that it seeks to preserve detail while smoothing nonimpulse noise, something that the "traditional" median filter does not do. The adaptive median filter also works in a rectangular window area S_{xy} . Unlike those filters, however, the adaptive median filter changes (increases) the size of S_{xy} during filter operation, depending on certain conditions. The output of the filter is a single value used to replace the value of the pixel at (x, y) , the particular point on which the window S_{xy} is centered at a given time.

Consider the following notation:

z_{\min} = minimum gray level value in S_{xy}

z_{\max} = maximum gray level value in S_{xy}

z_{med} = median of gray levels in S_{xy}

z_{xy} = gray level at coordinates (x, y)

S_{\max} = maximum allowed size of S_{xy} .

The adaptive median filtering algorithm works in two levels, denoted level A and level B, as follows:

Level A: $A1 = Z_{med} - Z_{min}$

$A2 = Z_{med} - Z_{max}$

If $A1 > 0$ AND $A2 < 0$, Go to level B

Else increase the window size

If window size $\leq S_{max}$ repeat level A

Else output z_{xy}

Level B: $B1 = Z_{xy} - Z_{min}$

$B2 = Z_{xy} - Z_{max}$

If $B1 > 0$ AND $B2 < 0$, output z_{xy}

Else output z_{med}

8. Explain a simple Image Formation Model.

An image is represented by two-dimensional functions of the form $f(x, y)$. The value or amplitude of f at spatial coordinates (x, y) is a positive scalar quantity whose physical meaning is determined by the source of the image. When an image is generated from a physical process, its values are proportional to energy radiated by a physical source (e.g., electromagnetic waves). As a consequence, $f(x, y)$ must be nonzero and finite; that is,

$$0 < f(x, y) < \infty \quad \dots (1)$$

The function $f(x, y)$ may be characterized by two components:

- A) The amount of source illumination incident on the scene being viewed.
- B) The amount of illumination reflected by the objects in the scene.

Appropriately, these are called the illumination and reflectance components and are denoted by $i(x, y)$ and $r(x, y)$, respectively. The two functions combine as a product to form $f(x, y)$.

$$f(x, y) = i(x, y) r(x, y) \quad \dots (2)$$

where

$$0 < i(x, y) < \infty \quad \dots (3)$$

and

$$0 < r(x, y) < 1 \quad \dots (4)$$

Equation (4) indicates that reflectance is bounded by 0 (total absorption) and 1 (total reflectance). The nature of $i(x, y)$ is determined by the illumination source, and $r(x, y)$ is determined by the characteristics of the imaged objects. It is noted that these expressions also are applicable to images formed via transmission of the illumination through a medium, such as a chest X-ray.

9. Write brief notes on inverse filtering.

The simplest approach to restoration is direct inverse filtering, where $F(u, v)$, the transform of the original image is computed simply by dividing the transform of the degraded image, $G(u, v)$, by the degradation function

$$\hat{F}(u, v) = \frac{G(u, v)}{H(u, v)}.$$

The divisions are between individual elements of the functions.

But $G(u, v)$ is given by

$$G(u, v) = F(u, v) + N(u, v)$$

Hence

$$\hat{F}(u, v) = F(u, v) + \frac{N(u, v)}{H(u, v)}.$$

It tells that even if the degradation function is known the undegraded image cannot be recovered [the inverse Fourier transform of $F(u, v)$] exactly because $N(u, v)$ is a random function whose Fourier transform is not known.

If the degradation has zero or very small values, then the ratio $N(u, v)/H(u, v)$ could easily dominate the estimate $F(u, v)$.

One approach to get around the zero or small-value problem is to limit the filter frequencies to values near the origin. $H(0, 0)$ is equal to the average value of $h(x, y)$ and that this

is usually the highest value of $H(u, v)$ in the frequency domain. Thus, by limiting the analysis to frequencies near the origin, the probability of encountering zero values is reduced.

10. Write about Noise Probability Density Functions.

The following are among the most common PDFs found in image processing applications.

Gaussian noise

Because of its mathematical tractability in both the spatial and frequency domains, Gaussian (also called normal) noise models are used frequently in practice. In fact, this tractability is so convenient that it often results in Gaussian models being used in situations in which they are marginally applicable at best.

The PDF of a Gaussian random variable, z , is given by

$$p(z) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(z-\mu)^2/2\sigma^2} \dots (1)$$

where z represents gray level, μ is the mean of average value of z , and a σ is its standard deviation. The standard deviation squared, σ^2 , is called the variance of z . A plot of this function is shown in Fig. 5.10. When z is described by Eq. (1), approximately 70% of its values will be in the range $[(\mu - \sigma), (\mu + \sigma)]$, and about 95% will be in the range $[(\mu - 2\sigma), (\mu + 2\sigma)]$.

Rayleigh noise

The PDF of Rayleigh noise is given by

$$p(z) = \begin{cases} \frac{2}{b} (z - a)e^{-(z-a)^2/b} & \text{for } z \geq a \\ 0 & \text{for } z < a. \end{cases}$$

The mean and variance of this density are given by

$$\mu = a + \sqrt{\Pi b/4}$$

$$\sigma^2 = b(4 - \Pi)/4$$

Figure 5.10 shows a plot of the Rayleigh density. Note the displacement from the origin and the fact that the basic shape of this density is skewed to the right. The Rayleigh density can be quite useful for approximating skewed histograms.

Erlang (Gamma) noise

The PDF of Erlang noise is given by

$$p(z) = \begin{cases} \frac{a^b z^{b-1}}{(b-1)!} e^{-az} & \text{for } z \geq 0 \\ 0 & \text{for } z < 0 \end{cases}$$

where the parameters are such that $a > 0$, b is a positive integer, and "!" indicates factorial. The mean and variance of this density are given by

$$\mu = b / a$$

$$\sigma^2 = b / a^2$$

Exponential noise

The PDF of exponential noise is given by

$$p(z) = \begin{cases} ae^{-az} & \text{for } z \geq 0 \\ 0 & \text{for } z < 0 \end{cases}$$

The mean of this density function is given by

$$\mu = 1 / a$$

$$\sigma^2 = 1 / a^2$$

This PDF is a special case of the Erlang PDF, with $b = 1$.

Uniform noise

The PDF of uniform noise is given by

$$p(z) = \begin{cases} \frac{1}{b-a} & \text{if } a \leq z \leq b \\ 0 & \text{otherwise.} \end{cases}$$

The mean of this density function is given by

$$\mu = a + b / 2$$

$$\sigma^2 = (b - a)^2 / 12$$

Impulse (salt-and-pepper) noise

The PDF of (bipolar) impulse noise is given by

$$p(z) = \begin{cases} P_a & \text{for } z = a \\ P_b & \text{for } z = b \\ 0 & \text{otherwise} \end{cases}$$

If $b > a$, gray-level b will appear as a light dot in the image. Conversely, level a will appear like a dark dot. If either P_a or P_b is zero, the impulse noise is called unipolar. If neither probability is zero, and especially if they are approximately equal, impulse noise values will resemble salt-and-pepper granules randomly distributed over the image. For this reason, bipolar impulse noise also is called salt-and-pepper noise. Shot and spike noise also are terms used to refer to this type of noise.

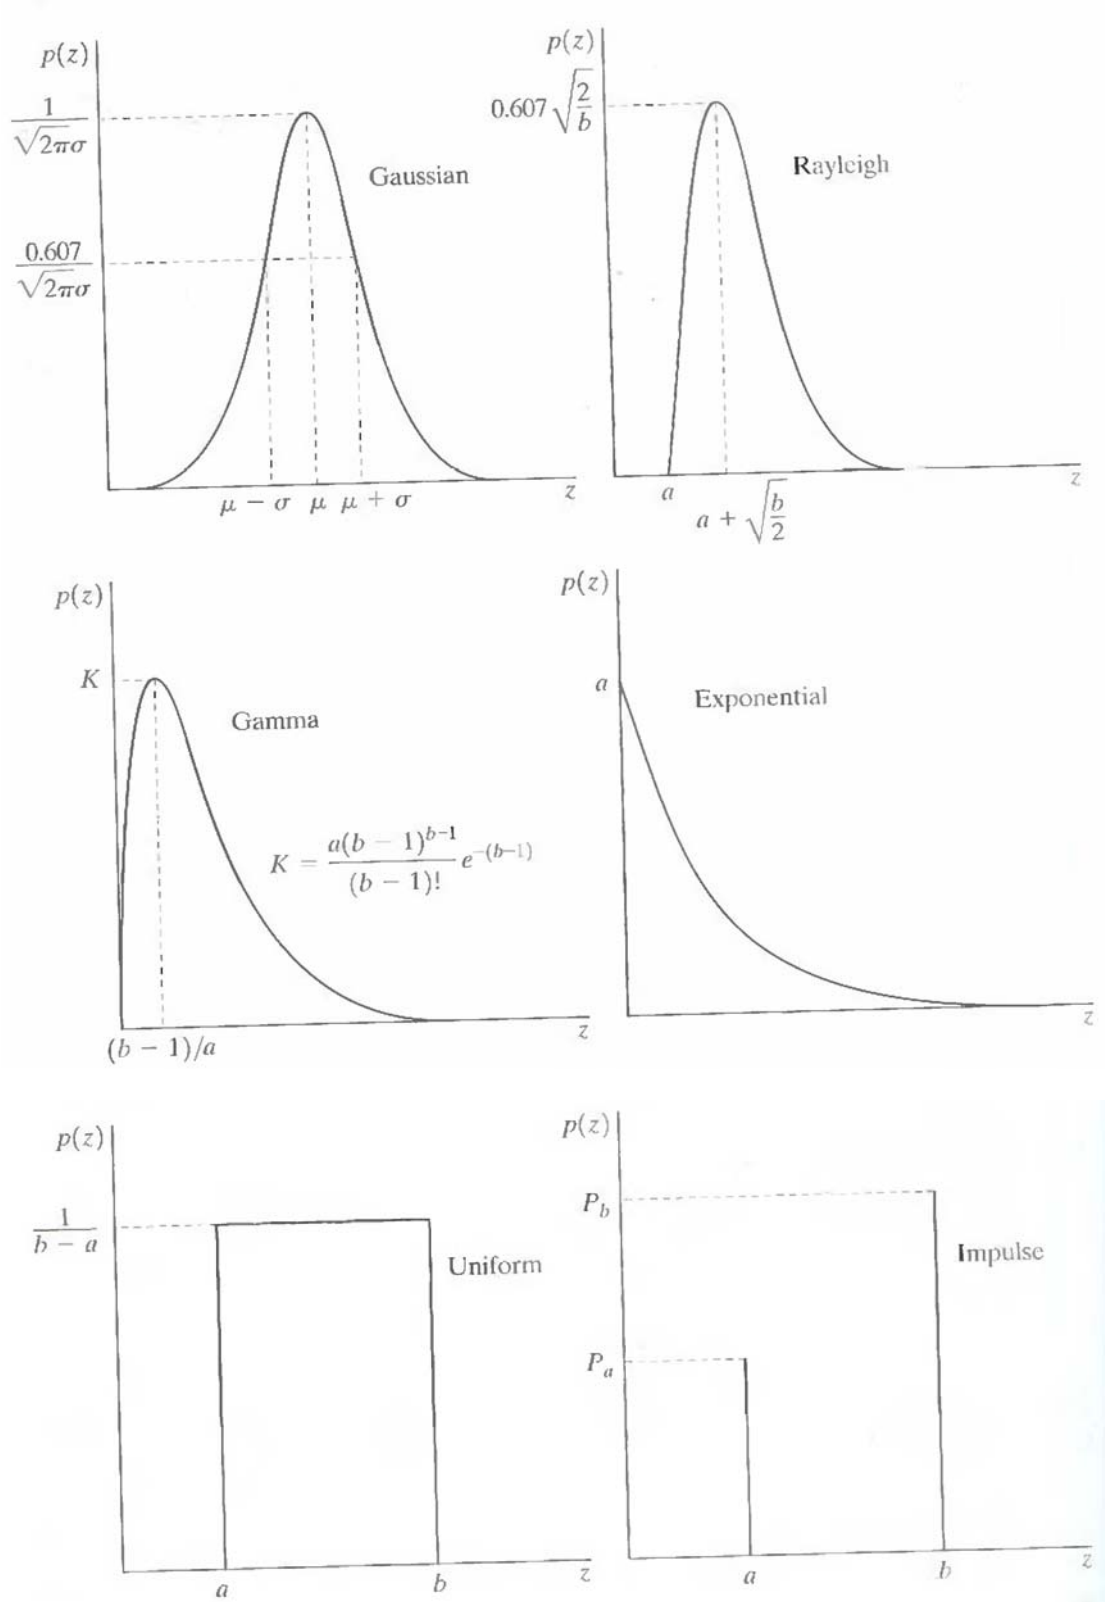


Fig.5.10 Some important probability density functions

11. Enumerate the differences between the image enhancement and image restoration.

(i) Image enhancement techniques are heuristic procedures designed to manipulate an image in order to take advantage of the psychophysical aspects of the human system. Whereas image restoration techniques are basically reconstruction techniques by which a degraded image is reconstructed by using some of the prior knowledge of the degradation phenomenon.

(ii) Image enhancement can be implemented by spatial and frequency domain technique, whereas image restoration can be implement by frequency domain and algebraic techniques.

(iii) The computational complexity for image enhancement is relatively less when compared to the computational complexity for irrrage restoration, since algebraic methods requires manipulation of large number of simultaneous equation. But, under some condition computational complexity can be reduced to the same level as that required by traditional frequency domain technique.

(iv) Image enhancement techniques are problem oriented, whereas image restoration techniques are general and are oriented towards modeling the degradation and applying the reverse process in order to reconstruct the original image.

(v) Masks are used in spatial domain methods for image enhancement, whereas masks are not used for image restoration techniques.

(vi) Contrast stretching is considered as image enhancement technique because it is based on the pleasing aspects of the review, whereas removal of image blur by applying a deblurring function is considered as a image restoration technique.

12. Explain about iterative nonlinear restoration using the Lucy–Richardson algorithm.

Lucy-Richardson algorithm is a nonlinear restoration method used to recover a latent image which is blurred by a Point Spread Function (psf). It is also known as Richardson-Lucy deconvolution.

With as the point spread function, the pixels in observed image are expressed as,

$$C_i = \sum_j P_{ij} u_j$$

Here,

u_j = Pixel value at location j in the image

c_i = Observed value at i^{th} pixel location

The L-R algorithm cannot be used in application in which the psf (P_{ij}) is dependent on one or more unknown variables.

The L-R algorithm is based on maximum-likelihood formulation, in this formulation Poisson statistics are used to model the image. If the likelihood of model is increased, then the result is an equation which satisfies when the following iteration converges.

$$\hat{f}_{k+1}(x, y) = \hat{f}_k(x, y) \left[h(-x, -y) \times \frac{g(x, y)}{h(x, y) \times \hat{f}_k(x, y)} \right]$$

Here,

f = Estimation of undegraded image.

The factor f which is present in the right side denominator leads to non-linearity. Since, the algorithm is a type of nonlinear restorations; hence it is stopped when satisfactory result is obtained.

The basic syntax of function deconvlucy with the L-R algorithm is implemented is given below.

fr = Deconvlucy (g, psf, NUMIT, DAMPAR, WEIGHT)

Here the parameters are,

g = Degraded image

f_r = Restored image

psf = Point spread function

NUMIT = Total number of iterations.

The remaining two parameters are,

DAMPAR

The DAMPAR parameter is a scalar parameter which is used to determine the deviation of resultant image with the degraded image (g). The pixels which get deviated from their original value within the DAMPAR, for these pixels iterations are cancelled so as to reduce noise generation and present essential image information.

WEIGHT

WEIGHT parameter gives a weight to each and every pixel. It is array of size similar to that of degraded image (g). In applications where a pixel leads to improper image is removed by assigning it to a weight as 0'. The pixels may also be given weights depending upon the flat-field correction, which is essential according to image array. Weights are used in applications such as blurring with specified psf. They are used to remove the pixels which are present at the boundary of the image and are blurred separately by psf.

If the array size of psf is $n \times n$ then the width of weight of border of zeroes being used is $\text{ceil}(n / 2)$