1. What is meant by Digital Image Processing? Explain how digital images can be represented?

An image may be defined as a two-dimensional function, f(x, y), where x and y are spatial (plane) coordinates, and the amplitude of f at any pair of coordinates (x, y) is called the intensity or gray level of the image at that point. When x, y, and the amplitude values of f are all finite, discrete quantities, we call the image a digital image. The field of digital image processing refers to processing digital images by means of a digital computer. Note that a digital image is composed of a finite number of elements, each of which has a particular location and value. These elements are referred to as picture elements, image elements, pels, and pixels. Pixel is the term most widely used to denote the elements of a digital image.

Vision is the most advanced of our senses, so it is not surprising that images play the single most important role in human perception. However, unlike humans, who are limited to the visual band of the electromagnetic (EM) spectrum, imaging machines cover almost the entire EM spectrum, ranging from gamma to radio waves. They can operate on images generated by sources that humans are not accustomed to associating with images. These include ultra-sound, electron microscopy, and computer-generated images. Thus, digital image processing encompasses a wide and varied field of applications. There is no general agreement among authors regarding where image processing stops and other related areas, such as image analysis and computer vision, start. Sometimes a distinction is made by defining image processing as a discipline in which both the input and output of a process are images. We believe this to be a limiting and somewhat artificial boundary. For example, under this definition, even the trivial task of computing the average intensity of an image (which yields a single number) would not be considered an image processing operation. On the other hand, there are fields such as computer vision whose ultimate goal is to use computers to emulate human vision, including learning and being able to make inferences and take actions based on visual inputs. This area itself is a branch of artificial intelligence (AI) whose objective is to emulate human intelligence. The field of AI is in its earliest stages of infancy in terms of development, with progress having been much slower than originally anticipated. The area of image analysis (also called image understanding) is in between image processing and computer vision.

There are no clear-cut boundaries in the continuum from image processing at one end to computer vision at the other. However, one useful paradigm is to consider three types of computerized processes in this continuum: low-, mid-, and high-level processes. Low-level processes involve primitive operations such as image preprocessing to reduce noise, contrast enhancement, and image sharpening. A low-level process is characterized by the fact that both its inputs and outputs are images. Mid-level processing on images involves tasks such as segmentation (partitioning an image into regions or objects), description of those objects to reduce them to a form suitable for computer processing, and classification (recognition) of individual objects. A mid-level process is characterized by the fact that its inputs generally are

images, but its outputs are attributes extracted from those images (e.g., edges, contours, and the identity of individual objects). Finally, higher-level processing involves "making sense" of an ensemble of recognized objects, as in image analysis, and, at the far end of the continuum, performing the cognitive functions normally associated with vision and, in addition, encompasses processes that extract attributes from images, up to and including the recognition of individual objects. As a simple illustration to clarify these concepts, consider the area of automated analysis of text. The processes of acquiring an image of the area containing the text, preprocessing that image, extracting (segmenting) the individual characters, describing the characters in a form suitable for computer processing, and recognizing those individual characters are in the scope of what we call digital image processing.

Representing Digital Images:

We will use two principal ways to represent digital images. Assume that an image $f(x, y)$ is sampled so that the resulting digital image has M rows and N columns. The values of the coordinates (x, y) now become discrete quantities. For notational clarity and convenience, we shall use integer values for these discrete coordinates. Thus, the values of the coordinates at the origin are $(x, y) = (0, 0)$. The next coordinate values along the first row of the image are represented as $(x, y) = (0, 1)$. It is important to keep in mind that the notation $(0, 1)$ is used to signify the second sample along the first row. It does not mean that these are the actual values of physical coordinates when the image was sampled. Figure 1 shows the coordinate convention used.
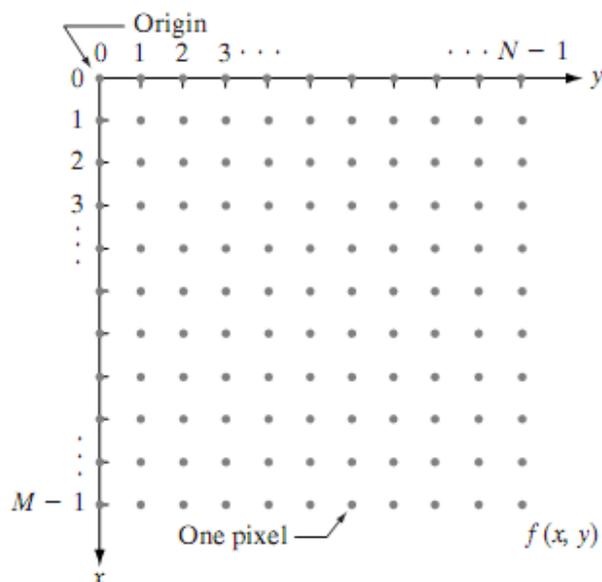

Fig 1 Coordinate convention used to represent digital images

The notation introduced in the preceding paragraph allows us to write the complete M*N digital image in the following compact matrix form:

$$
f(x, y) = \begin{bmatrix} f(0,0) & f(0,1) & \cdots & f(0, N-1) \\ f(1,0) & f(1,1) & \cdots & f(1, N-1) \\ \vdots & \vdots & & \vdots \\ f(M-1,0) & f(M-1,1) & \cdots & f(M-1, N-1) \end{bmatrix}.
$$

The right side of this equation is by definition a digital image. Each element of this matrix array is called an image element, picture element, pixel, or pel.

2. What are the fundamental steps in Digital Image Processing?

Fundamental Steps in Digital Image Processing:

Image acquisition is the first process shown in Fig.2. Note that acquisition could be as simple as being given an image that is already in digital form. Generally, the image acquisition stage involves preprocessing, such as scaling.

Image enhancement is among the simplest and most appealing areas of digital image processing. Basically, the idea behind enhancement techniques is to bring out detail that is obscured, or simply to highlight certain features of interest in an image. A familiar example of enhancement is when we increase the contrast of an image because "it looks better." It is important to keep in mind that enhancement is a very subjective area of image processing.

Image restoration is an area that also deals with improving the appearance of an image. However, unlike enhancement, which is subjective, image restoration is objective, in the sense that restoration techniques tend to be based on mathematical or probabilistic models of image degradation. Enhancement, on the other hand, is based on human subjective preferences regarding what constitutes a "good" enhancement result.

Color image processing is an area that has been gaining in importance because of the significant increase in the use of digital images over the Internet.
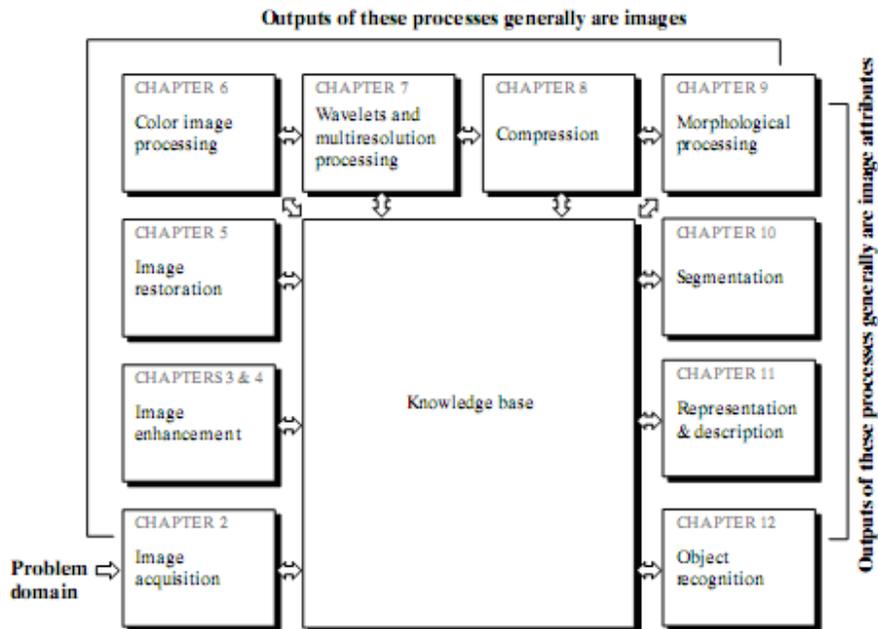
Fig.2. Fundamental steps in Digital Image Processing

Wavelets are the foundation for representing images in various degrees of resolution. Compression, as the name implies, deals with techniques for reducing the storage required to save an image, or the bandwidth required to transmit it. Although storage technology has improved significantly over the past decade, the same cannot be said for transmission capacity. This is true particularly in uses of the Internet, which are characterized by significant pictorial content. Image compression is familiar (perhaps inadvertently) to most users of computers in the form of image file extensions, such as the jpg file extension used in the JPEG (Joint Photographic Experts Group) image compression standard.

Morphological processing deals with tools for extracting image components that are useful in the representation and description of shape.

Segmentation procedures partition an image into its constituent parts or objects. In general, autonomous segmentation is one of the most difficult tasks in digital image processing. A rugged segmentation procedure brings the process a long way toward successful solution of imaging problems that require objects to be identified individually. On the other hand, weak or erratic segmentation algorithms almost always guarantee eventual failure. In general, the more accurate the segmentation, the more likely recognition is to succeed.

Representation and description almost always follow the output of a segmentation stage, which usually is raw pixel data, constituting either the boundary of a region (i.e., the set of pixels separating one image region from another) or all the points in the region itself. In either case, converting the data to a form suitable for computer processing is necessary. The first decision that must be made is whether the data should be represented as a boundary or as a complete region. Boundary representation is appropriate when the focus is on external shape characteristics, such as corners and inflections. Regional representation is appropriate when the focus is on internal properties, such as texture or skeletal shape. In some applications, these representations complement each other. Choosing a representation is only part of the solution for transforming raw data into a form suitable for subsequent computer processing. A method must also be specified for describing the data so that features of interest are highlighted. Description, also called feature selection, deals with extracting attributes that result in some quantitative information of interest or are basic for differentiating one class of objects from another.

Recognition is the process that assigns a label (e.g., "vehicle") to an object based on its descriptors. We conclude our coverage of digital image processing with the development of methods for recognition of individual objects.

3. What are the components of an Image Processing System?

Components of an Image Processing System:

As recently as the mid-1980s, numerous models of image processing systems being sold throughout the world were rather substantial peripheral devices that attached to equally substantial host computers. Late in the 1980s and early in the 1990s, the market shifted to image processing hardware in the form of single boards designed to be compatible with industry standard buses and to fit into engineering workstation cabinets and personal computers. In addition to lowering costs, this market shift also served as a catalyst for a significant number of new companies whose specialty is the development of software written specifically for image processing.

Although large-scale image processing systems still are being sold for massive imaging applications, such as processing of satellite images, the trend continues toward miniaturizing and blending of general-purpose small computers with specialized image processing hardware. Figure 3 shows the basic components comprising a typical general-purpose system used for digital image processing. The function of each component is discussed in the following paragraphs, starting with image sensing.

With reference to sensing, two elements are required to acquire digital images. The first is a physical device that is sensitive to the energy radiated by the object we wish to image. The second, called a digitizer, is a device for converting the output of the physical sensing device into

digital form. For instance, in a digital video camera, the sensors produce an electrical output proportional to light intensity. The digitizer converts these outputs to digital data.

Specialized image processing hardware usually consists of the digitizer just mentioned, plus hardware that performs other primitive operations, such as an arithmetic logic unit (ALU), which performs arithmetic and logical operations in parallel on entire images. One example of how an ALU is used is in averaging images as quickly as they are digitized, for the purpose of noise reduction. This type of hardware sometimes is called a front-end subsystem, and its most distinguishing characteristic is speed. In other words, this unit performs functions that require fast data throughputs (e.g., digitizing and averaging video images at 30 framess) that the typical main computer cannot handle.
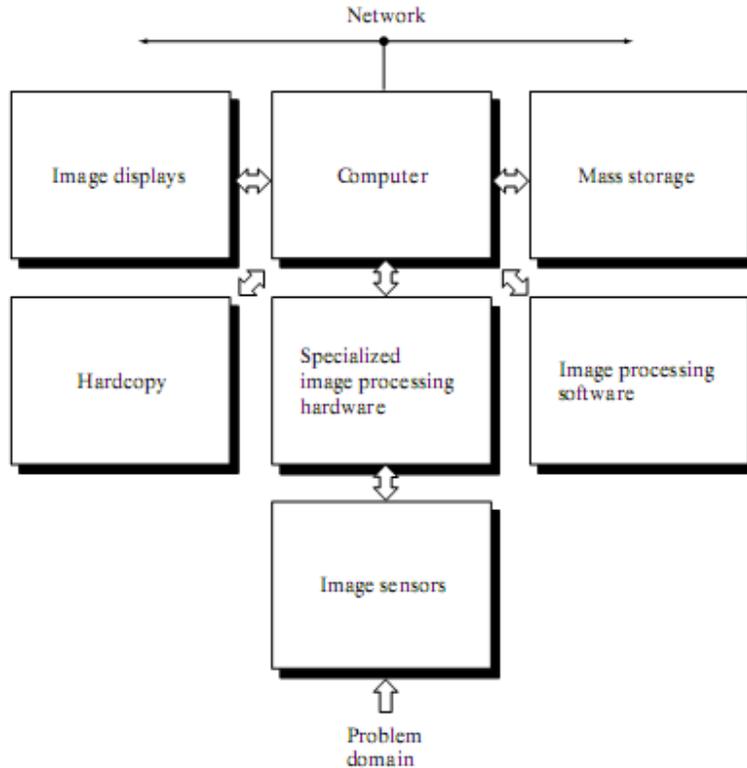
Fig.3. Components of a general purpose Image Processing System

The computer in an image processing system is a general-purpose computer and can range from a PC to a supercomputer. In dedicated applications, some times specially designed computers are used to achieve a required level of performance, but our interest here is on general-purpose

image processing systems. In these systems, almost any well-equipped PC-type machine is suitable for offline image processing tasks.

Software for image processing consists of specialized modules that perform specific tasks. A well-designed package also includes the capability for the user to write code that, as a minimum, utilizes the specialized modules. More sophisticated software packages allow the integration of those modules and general-purpose software commands from at least one computer language.

Mass storage capability is a must in image processing applications. An image of size 1024*1024 pixels, in which the intensity of each pixel is an 8-bit quantity, requires one megabyte of storage space if the image is not compressed. When dealing with thousands, or even millions, of images, providing adequate storage in an image processing system can be a challenge. Digital storage for image processing applications falls into three principal categories: (1) short-term storage for use during processing, (2) on-line storage for relatively fast re-call, and (3) archival storage, characterized by infrequent access. Storage is measured in bytes (eight bits), Kbytes (one thousand bytes), Mbytes (one million bytes), Gbytes (meaning giga, or one billion, bytes), and Tbytes (meaning tera, or one trillion, bytes). One method of providing short-term storage is computer memory. Another is by specialized boards, called frame buffers, that store one or more images and can be accessed rapidly, usually at video rates (e.g., at 30 complete images per second).The latter method allows virtually instantaneous image zoom, as well as scroll (vertical shifts) and pan (horizontal shifts). Frame buffers usually are housed in the specialized image processing hardware unit shown in Fig.3.Online storage generally takes the form of magnetic disks or optical-media storage. The key factor characterizing on-line storage is frequent access to the stored data. Finally, archival storage is characterized by massive storage requirements but infrequent need for access. Magnetic tapes and optical disks housed in "jukeboxes" are the usual media for archival applications.

Image displays in use today are mainly color (preferably flat screen) TV monitors. Monitors are driven by the outputs of image and graphics display cards that are an integral part of the computer system. Seldom are there requirements for image display applications that cannot be met by display cards available commercially as part of the computer system. In some cases, it is necessary to have stereo displays, and these are implemented in the form of headgear containing two small displays embedded in goggles worn by the user.

Hardcopy devices for recording images include laser printers, film cameras, heat-sensitive devices, inkjet units, and digital units, such as optical and CD-ROM disks. Film provides the highest possible resolution, but paper is the obvious medium of choice for written material. For presentations, images are displayed on film transparencies or in a digital medium if image projection equipment is used. The latter approach is gaining acceptance as the standard for image presentations.

Networking is almost a default function in any computer system in use today. Because of the large amount of data inherent in image processing applications, the key consideration in image transmission is bandwidth. In dedicated networks, this typically is not a problem, but communications with remote sites via the Internet are not always as efficient. Fortunately, this situation is improving quickly as a result of optical fiber and other broadband technologies.

4. Explain about elements of visual perception.

Elements of Visual Perception:

Although the digital image processing field is built on a foundation of mathematical and probabilistic formulations, human intuition and analysis play a central role in the choice of one technique versus another, and this choice often is made based on subjective, visual judgments.

(1) Structure of the Human Eye:

Figure 4.1 shows a simplified horizontal cross section of the human eye. The eye is nearly a sphere, with an average diameter of approximately 20 mm. Three membranes enclose the eye: the cornea and sclera outer cover; the choroid; and the retina. The cornea is a tough, transparent tissue that covers the anterior surface of the eye. Continuous with the cornea, the sclera is an opaque membrane that encloses the remainder of the optic globe. The choroid lies directly below the sclera. This membrane contains a network of blood vessels that serve as the major source of nutrition to the eye. Even superficial injury to the choroid, often not deemed serious, can lead to severe eye damage as a result of inflammation that restricts blood flow. The choroid coat is heavily pigmented and hence helps to reduce the amount of extraneous light entering the eye and the backscatter within the optical globe. At its anterior extreme, the choroid is divided into the ciliary body and the iris diaphragm. The latter contracts or expands to control the amount of light that enters the eye. The central opening of the iris (the pupil) varies in diameter from approximately 2 to 8 mm. The front of the iris contains the visible pigment of the eye, whereas the back contains a black pigment.

The lens is made up of concentric layers of fibrous cells and is suspended by fibers that attach to the ciliary body. It contains 60 to 70%water, about 6%fat, and more protein than any other tissue in the eye. The lens is colored by a slightly yellow pigmentation that increases with age. In extreme cases, excessive clouding of the lens, caused by the affliction commonly referred to as cataracts, can lead to poor color discrimination and loss of clear vision. The lens absorbs approximately 8% of the visible light spectrum, with relatively higher absorption at shorter wavelengths. Both infrared and ultraviolet light are absorbed appreciably by proteins within the lens structure and, in excessive amounts, can damage the eye.
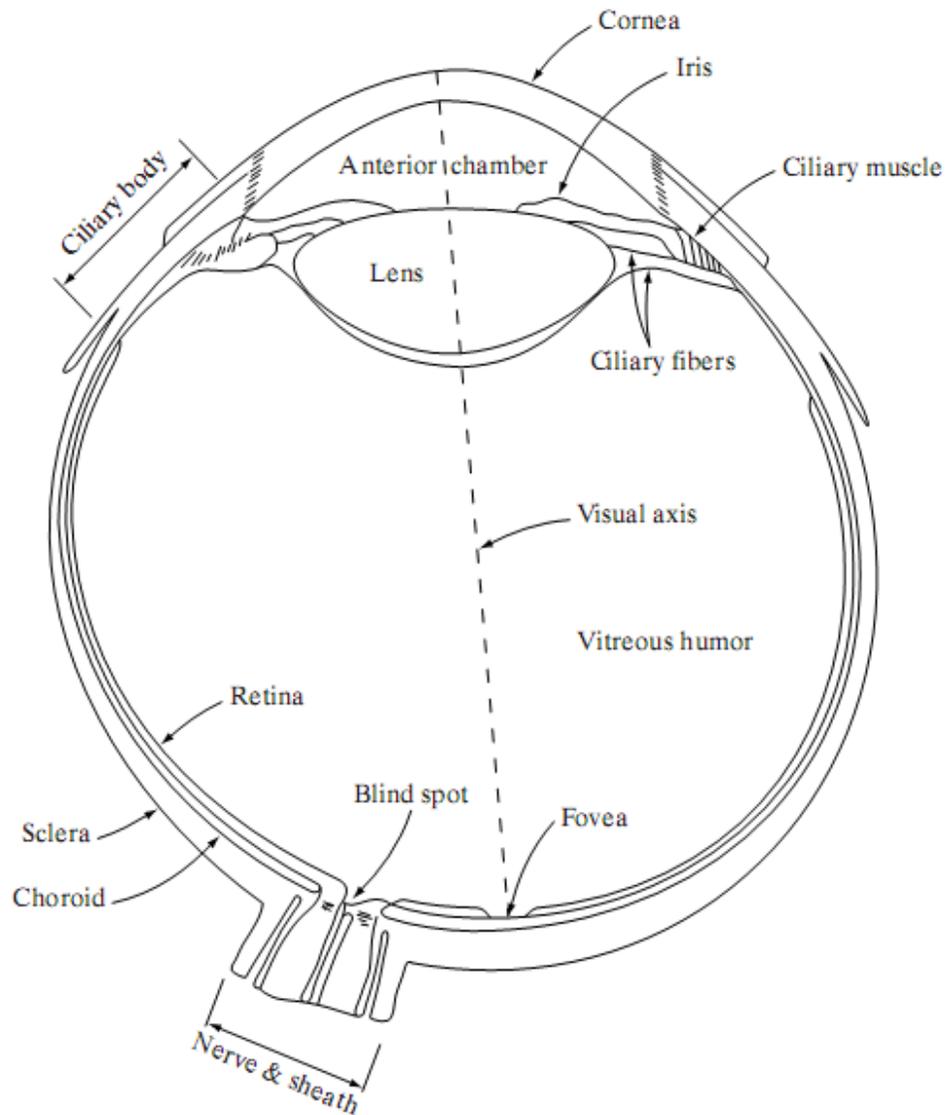
Fig.4.1 Simplified diagram of a cross section of the human eye.

The innermost membrane of the eye is the retina, which lines the inside of the wall's entire posterior portion. When the eye is properly focused, light from an object outside the eye is imaged on the retina. Pattern vision is afforded by the distribution of discrete light receptors over the surface of the retina. There are two classes of receptors: cones and rods. The cones in each eye number between 6 and 7 million. They are located primarily in the central portion of the

retina, called the fovea, and are highly sensitive to color. Humans can resolve fine details with these cones largely because each one is connected to its own nerve end. Muscles controlling the eye rotate the eyeball until the image of an object of interest falls on the fovea. Cone vision is called photopic or bright-light vision. The number of rods is much larger: Some 75 to 150 million are distributed over the retinal surface. The larger area of distribution and the fact that several rods are connected to a single nerve end reduce the amount of detail discernible by these receptors. Rods serve to give a general, overall picture of the field of view. They are not involved in color vision and are sensitive to low levels of illumination. For example, objects that appear brightly colored in daylight when seen by moonlight appear as colorless forms because only the rods are stimulated. This phenomenon is known as scotopic or dim-light vision.

(2) Image Formation in the Eye:

The principal difference between the lens of the eye and an ordinary optical lens is that the former is flexible. As illustrated in Fig. 4.1, the radius of curvature of the anterior surface of the lens is greater than the radius of its posterior surface. The shape of the lens is controlled by tension in the fibers of the ciliary body. To focus on distant objects, the controlling muscles cause the lens to be relatively flattened. Similarly, these muscles allow the lens to become thicker in order to focus on objects near the eye. The distance between the center of the lens and the retina (called the focal length) varies from approximately 17 mm to about 14 mm, as the refractive power of the lens increases from its minimum to its maximum. When the eye
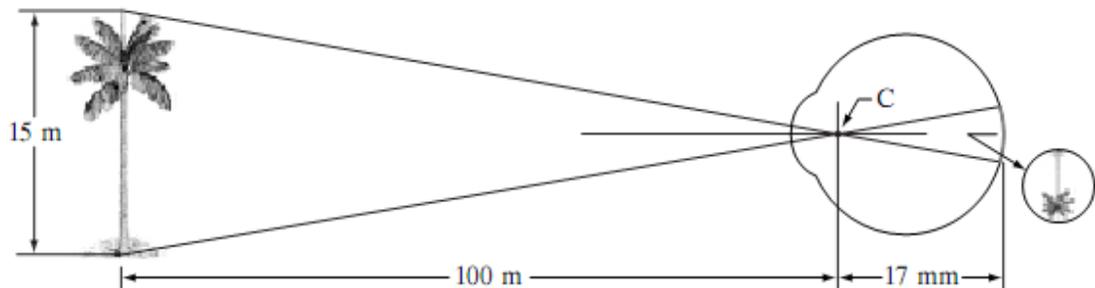
Fig.4.2. Graphical representation of the eye looking at a palm tree Point C is the optical center of the lens.

focuses on an object farther away than about 3 m, the lens exhibits its lowest refractive power. When the eye focuses on a nearby object, the lens is most strongly refractive. This information

makes it easy to calculate the size of the retinal image of any object. In Fig. 4.2, for example, the observer is looking at a tree 15 m high at a distance of 100 m. If h is the height in mm of that object in the retinal image, the geometry of Fig.4.2 yields 15/100=h/17 or h=2.55mm. The retinal image is reflected primarily in the area of the fovea. Perception then takes place by the relative excitation of light receptors, which transform radiant energy into electrical impulses that are ultimately decoded by the brain.

(3)Brightness Adaptation and Discrimination:

Because digital images are displayed as a discrete set of intensities, the eye's ability to discriminate between different intensity levels is an important consideration in presenting image-processing results. The range of light intensity levels to which the human visual system can adapt is enormous—on the order of 10^{10}—from the scotopic threshold to the glare limit. Experimental evidence indicates that subjective brightness (intensity as perceived by the human visual system) is a logarithmic function of the light intensity incident on the eye. Figure 4.3, a plot of light intensity versus subjective brightness, illustrates this characteristic. The long solid curve represents the range of intensities to which the visual system can adapt. In photopic vision alone, the range is about 10^6. The transition from scotopic to photopic vision is gradual over the approximate range from 0.001 to 0.1 millilambert (−3 to −1 mL in the log scale), as the double branches of the adaptation curve in this range show.
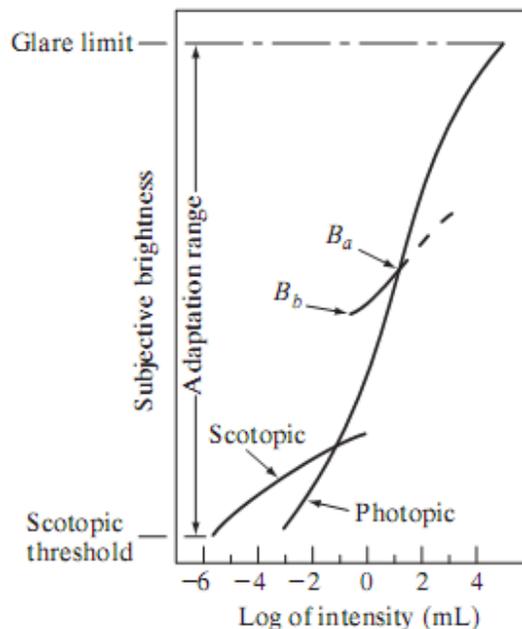
Fig.4.3. Range of Subjective brightness sensations showing a particular adaptation level.

The essential point in interpreting the impressive dynamic range depicted in Fig.4.3 is that the visual system cannot operate over such a range simultaneously. Rather, it accomplishes this large variation by changes in its overall sensitivity, a phenomenon known as brightness adaptation. The total range of distinct intensity levels it can discriminate simultaneously is rather small when compared with the total adaptation range. For any given set of conditions, the current sensitivity level of the visual system is called the brightness adaptation level, which may correspond, for example, to brightness Ba in Fig. 4.3. The short intersecting curve represents the range of subjective brightness that the eye can perceive when adapted to this level. This range is rather restricted, having a level Bb at and below which all stimuli are perceived as indistinguishable blacks. The upper (dashed) portion of the curve is not actually restricted but, if extended too far, loses its meaning because much higher intensities would simply raise the adaptation level higher than Ba.

5. Explain the process of image acquisition.

Image Sensing and Acquisition:

The types of images in which we are interested are generated by the combination of an "illumination" source and the reflection or absorption of energy from that source by the elements of the "scene" being imaged. We enclose illumination and scene in quotes to emphasize the fact that they are considerably more general than the familiar situation in which a visible light source illuminates a common everyday 3-D (three-dimensional) scene. For example, the illumination may originate from a source of electromagnetic energy such as radar, infrared, or X-ray energy. But, as noted earlier, it could originate from less traditional sources, such as ultrasound or even a computer-generated illumination pattern.

Similarly, the scene elements could be familiar objects, but they can just as easily be molecules, buried rock formations, or a human brain. We could even image a source, such as acquiring images of the sun. Depending on the nature of the source, illumination energy is reflected from, or transmitted through, objects. An example in the first category is light reflected from a planar surface. An example in the second category is when X-rays pass through a patient's body for the purpose of generating a diagnostic X-ray film. In some applications, the reflected or transmitted energy is focused onto a photo converter (e.g., a phosphor screen), which converts the energy into visible light. Electron microscopy and some applications of gamma imaging use this approach.

Figure 5.1 shows the three principal sensor arrangements used to transform illumination energy into digital images. The idea is simple: Incoming energy is transformed into a voltage by the combination of input electrical power and sensor material that is responsive to the particular type of energy being detected. The output voltage waveform is the response of the sensor(s), and a digital quantity is obtained from each sensor by digitizing its response.
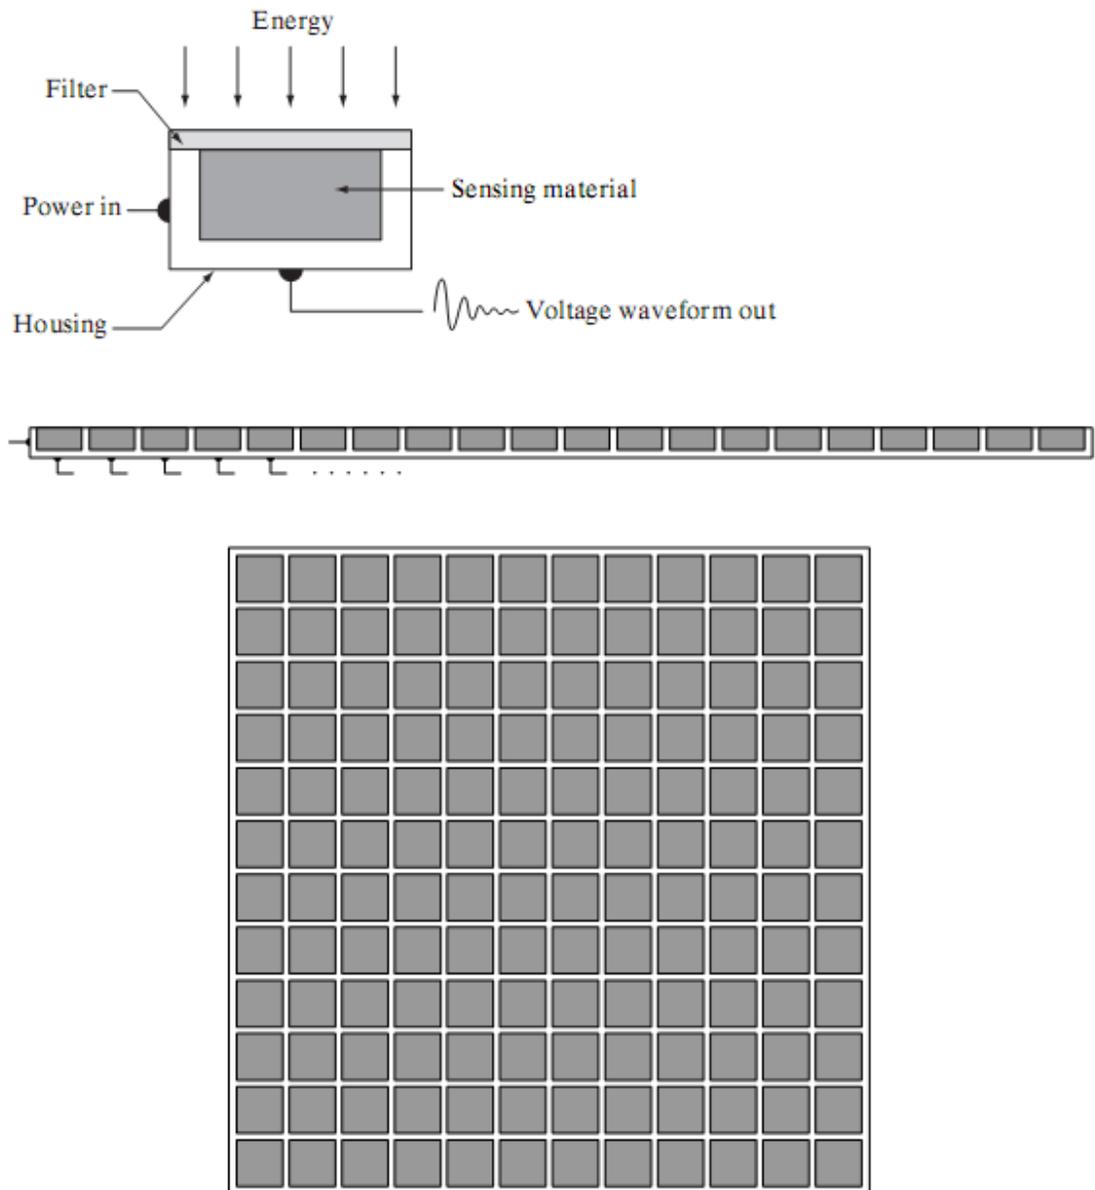
Fig.5.1 (a) Single imaging Sensor (b) Line sensor (c) Array sensor

(1)Image Acquisition Using a Single Sensor:

Figure 5.1 (a) shows the components of a single sensor. Perhaps the most familiar sensor of this type is the photodiode, which is constructed of silicon materials and whose output voltage waveform is proportional to light. The use of a filter in front of a sensor improves selectivity. For example, a green (pass) filter in front of a light sensor favors light in the green band of the color

spectrum. As a consequence, the sensor output will be stronger for green light than for other components in the visible spectrum.

In order to generate a 2-D image using a single sensor, there has to be relative displacements in both the x- and y-directions between the sensor and the area to be imaged. Figure 5.2 shows an arrangement used in high-precision scanning, where a film negative is mounted onto a drum whose mechanical rotation provides displacement in one dimension. The single sensor is mounted on a lead screw that provides motion in the perpendicular direction. Since mechanical motion can be controlled with high precision, this method is an inexpensive (but slow) way to obtain high-resolution images. Other similar mechanical arrangements use a flat bed, with the sensor moving in two linear directions. These types of mechanical digitizers sometimes are referred to as microdensitometers.
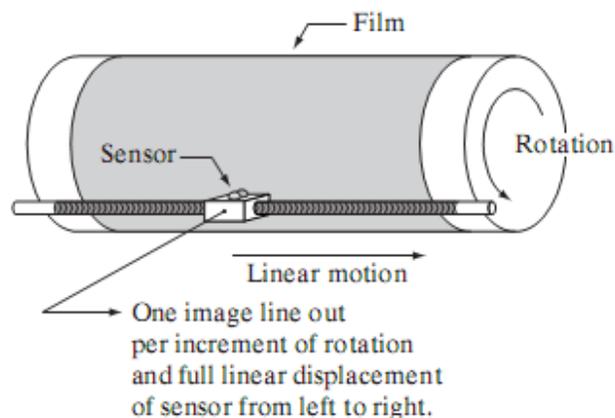
Fig.5.2. Combining a single sensor with motion to generate a 2-D image

(2) Image Acquisition Using Sensor Strips:

A geometry that is used much more frequently than single sensors consists of an in-line arrangement of sensors in the form of a sensor strip, as Fig. 5.1 (b) shows. The strip provides imaging elements in one direction. Motion perpendicular to the strip provides imaging in the other direction, as shown in Fig. 5.3 (a).This is the type of arrangement used in most flat bed scanners. Sensing devices with 4000 or more in-line sensors are possible. In-line sensors are used routinely in airborne imaging applications, in which the imaging system is mounted on an aircraft that flies at a constant altitude and speed over the geographical area to be imaged. One-dimensional imaging sensor strips that respond to various bands of the electromagnetic spectrum are mounted perpendicular to the direction of flight. The imaging strip gives one line of an image

at a time, and the motion of the strip completes the other dimension of a two-dimensional image. Lenses or other focusing schemes are used to project the area to be scanned onto the sensors.

Sensor strips mounted in a ring configuration are used in medical and industrial imaging to obtain cross-sectional ("slice") images of 3-D objects, as Fig. 5.3 (b) shows. A rotating X-ray source provides illumination and the portion of the sensors opposite the source collect the X-ray energy that pass through the object (the sensors obviously have to be sensitive to X-ray energy).This is the basis for medical and industrial computerized axial tomography (CAT). It is important to note that the output of the sensors must be processed by reconstruction algorithms whose objective is to transform the sensed data into meaningful cross-sectional images.

In other words, images are not obtained directly from the sensors by motion alone; they require extensive processing. A 3-D digital volume consisting of stacked images is generated as the object is moved in a direction perpendicular to the sensor ring. Other modalities of imaging based on the CAT principle include magnetic resonance imaging (MRI) and positron emission tomography (PET).The illumination sources, sensors, and types of images are different, but conceptually they are very similar to the basic imaging approach shown in Fig. 5.3 (b).
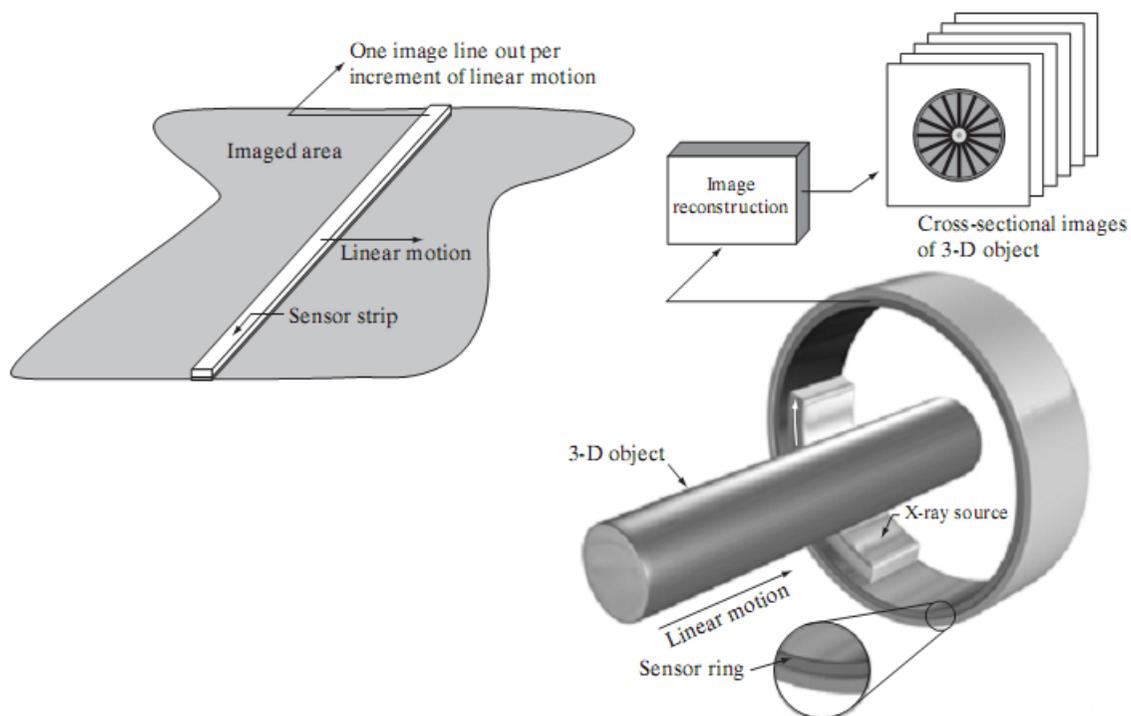
Fig.5.3 (a) Image acquisition using a linear sensor strip (b) Image acquisition using a circular sensor strip.

(3) Image Acquisition Using Sensor Arrays:

Figure 5.1 (c) shows individual sensors arranged in the form of a 2-D array. Numerous electromagnetic and some ultrasonic sensing devices frequently are arranged in an array format. This is also the predominant arrangement found in digital cameras. A typical sensor for these cameras is a CCD array, which can be manufactured with a broad range of sensing properties and can be packaged in rugged arrays of 4000 * 4000 elements or more. CCD sensors are used widely in digital cameras and other light sensing instruments. The response of each sensor is proportional to the integral of the light energy projected onto the surface of the sensor, a property that is used in astronomical and other applications requiring low noise images. Noise reduction is achieved by letting the sensor integrate the input light signal over minutes or even hours. Since the sensor array shown in Fig. 5.4 (c) is two dimensional, its key advantage is that a complete image can be obtained by focusing the energy pattern onto the surface of the array. The principal manner in which array sensors are used is shown in Fig.5.4. This figure shows the energy from an illumination source being reflected from a scene element, but, as mentioned at the beginning of this section, the energy also could be transmitted through the scene elements. The first function performed by the imaging system shown in Fig.5.4 (c) is to collect the incoming energy and focus it onto an image plane. If the illumination is light, the front end of the imaging system is a lens, which projects the viewed scene onto the lens focal plane, as Fig. 2.15(d) shows. The sensor array, which is coincident with the focal plane, produces outputs proportional to the integral of the light received at each sensor. Digital and analog circuitry sweep these outputs and converts them to a video signal, which is then digitized by another section of the imaging system. The output is a digital image, as shown diagrammatically in Fig. 5.4 (e).

Illumination (energy) source

Output (digitized) image

Imaging system

(Internal) image plane

Scene element

a
b c d e
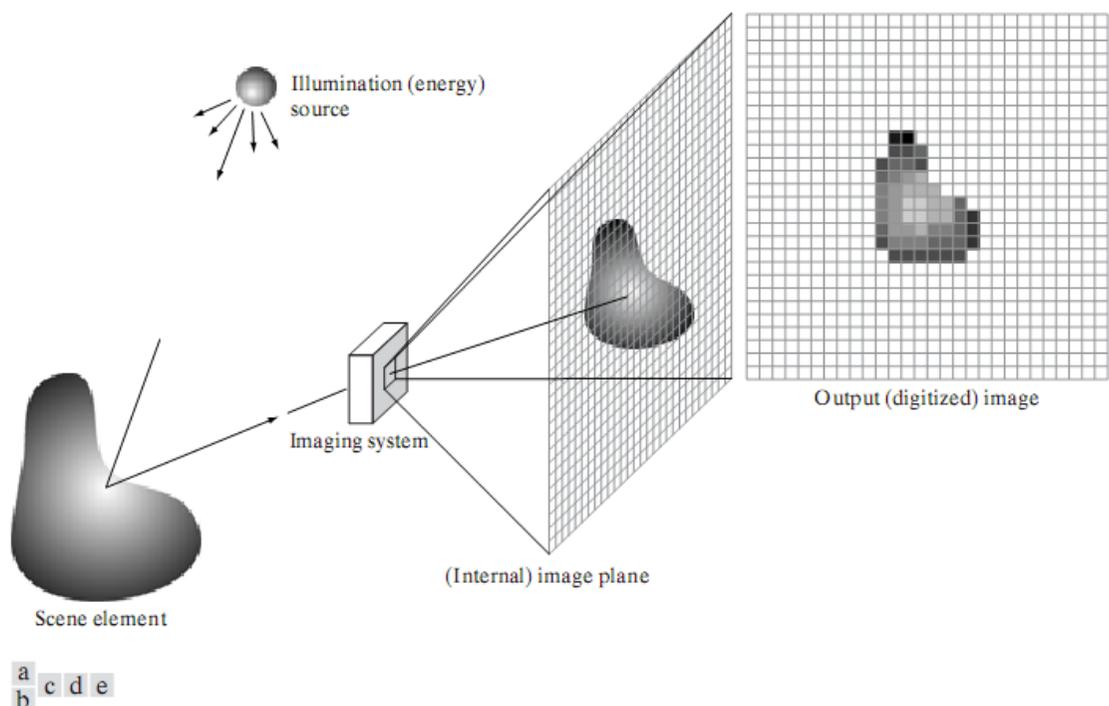
Fig.5.4 An example of the digital image acquisition process (a) Energy ("illumination") source (b) An element of a scene (c) Imaging system (d) Projection of the scene onto the image plane (e) Digitized image

6. Explain about image sampling and quantization process.

Image Sampling and Quantization:

The output of most sensors is a continuous voltage waveform whose amplitude and spatial behavior are related to the physical phenomenon being sensed. To create a digital image, we need to convert the continuous sensed data into digital form. This involves two processes: sampling and quantization.

Basic Concepts in Sampling and Quantization:

The basic idea behind sampling and quantization is illustrated in Fig.6.1. Figure 6.1(a) shows a continuous image, f(x, y), that we want to convert to digital form. An image may be continuous with respect to the x- and y-coordinates, and also in amplitude. To convert it to digital form, we have to sample the function in both coordinates and in amplitude. Digitizing the coordinate values is called sampling. Digitizing the amplitude values is called quantization.

The one-dimensional function shown in Fig.6.1 (b) is a plot of amplitude (gray level) values of the continuous image along the line segment AB in Fig. 6.1(a).The random variations are due to image noise. To sample this function, we take equally spaced samples along line AB, as shown in Fig.6.1 (c).The location of each sample is given by a vertical tick mark in the bottom part of the figure. The samples are shown as small white squares superimposed on the function. The set of these discrete locations gives the sampled function. However, the values of the samples still span (vertically) a continuous range of gray-level values. In order to form a digital function, the gray-level values also must be converted (quantized) into discrete quantities. The right side of Fig. 6.1 (c) shows the gray-level scale divided into eight discrete levels, ranging from black to white. The vertical tick marks indicate the specific value assigned to each of the eight gray levels. The continuous gray levels are quantized simply by assigning one of the eight discrete gray levels to each sample. The assignment is made depending on the vertical proximity of a sample to a vertical tick mark. The digital samples resulting from both sampling and quantization are shown in Fig.6.1 (d). Starting at the top of the image and carrying out this procedure line by line produces a two-dimensional digital image.

Sampling in the manner just described assumes that we have a continuous image in both coordinate directions as well as in amplitude. In practice, the method of sampling is determined by the sensor arrangement used to generate the image. When an image is generated by a single

sensing element combined with mechanical motion, as in Fig. 2.13, the output of the sensor is quantized in the manner described above. However, sampling is accomplished by selecting the number of individual mechanical increments at which we activate the sensor to collect data. Mechanical motion can be made very exact so, in principle; there is almost no limit as to how fine we can sample an image. However, practical limits are established by imperfections in the optics used to focus on the
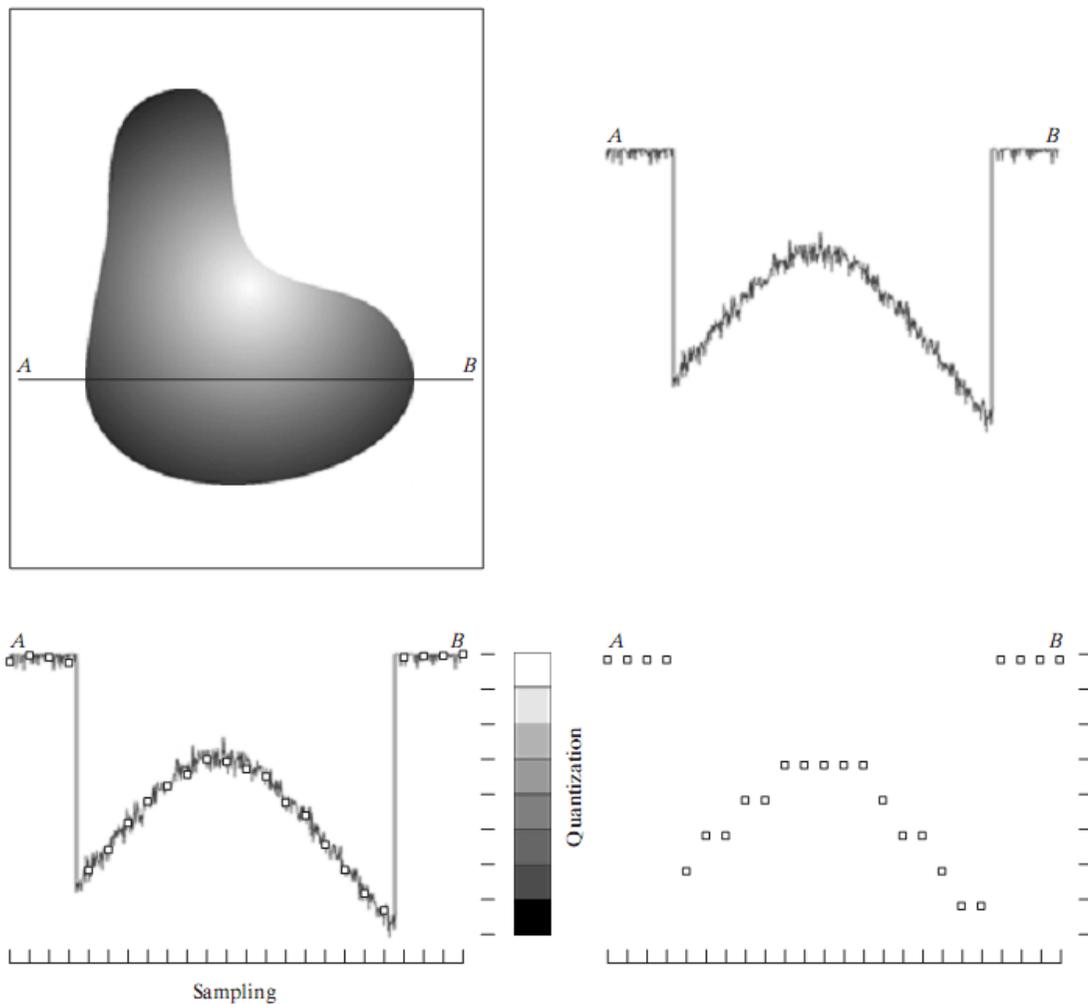
Fig.6.1. Generating a digital image (a) Continuous image (b) A scan line from A to Bin the continuous image, used to illustrate the concepts of sampling and quantization (c) Sampling and quantization. (d) Digital scan line

sensor an illumination spot that is inconsistent with the fine resolution achievable with mechanical displacements. When a sensing strip is used for image acquisition, the number of sensors in the strip establishes the sampling limitations in one image direction. Mechanical motion in the other direction can be controlled more accurately, but it makes little sense to try to achieve sampling density in one direction that exceeds the sampling limits established by the number of sensors in the other. Quantization of the sensor outputs completes the process of generating a digital image.

When a sensing array is used for image acquisition, there is no motion and the number of sensors in the array establishes the limits of sampling in both directions. Figure 6.2 illustrates this concept. Figure 6.2 (a) shows a continuous image projected onto the plane of an array sensor. Figure 6.2 (b) shows the image after sampling and quantization. Clearly, the quality of a digital image is determined to a large degree by the number of samples and discrete gray levels used in sampling and quantization.

a b

Fig.6.2. (a) Continuos image projected onto a sensor array (b) Result of image sampling and quantization.

7. Define spatial and gray level resolution. Explain about isopreference curves.

Spatial and Gray-Level Resolution:

Sampling is the principal factor determining the spatial resolution of an image. Basically, spatial resolution is the smallest discernible detail in an image. Suppose that we construct a chart with vertical lines of width W, with the space between the lines also having width W.A line pair consists of one such line and its adjacent space. Thus, the width of a line pair is 2W, and there are 1/2Wline pairs per unit distance. A widely used definition of resolution is simply the smallest number of discernible line pairs per unit distance; for example, 100 line pairs per millimeter. Gray-level resolution similarly refers to the smallest discernible change in gray level. We have considerable discretion regarding the number of samples used to generate a digital image, but this is not true for the number of gray levels. Due to hardware considerations, the number of gray levels is usually an integer power of 2.

The most common number is 8 bits, with 16 bits being used in some applications where enhancement of specific gray-level ranges is necessary. Sometimes we find systems that can digitize the gray levels of an image with 10 or 12 bit of accuracy, but these are the exception rather than the rule. When an actual measure of physical resolution relating pixels and the level of detail they resolve in the original scene are not necessary, it is not uncommon to refer to an L-level digital image of size M*N as having a spatial resolution of M*N pixels and a gray-level resolution of L levels.
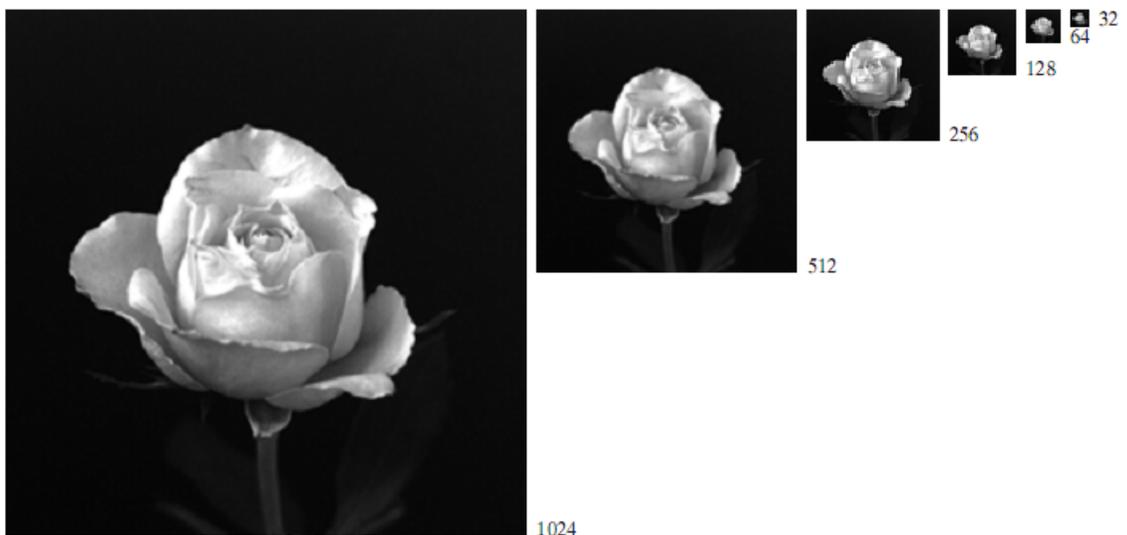
Fig.7.1. A 1024*1024, 8-bit image subsampled down to size 32*32 pixels The number of allowable gray levels was kept at 256.

The subsampling was accomplished by deleting the appropriate number of rows and columns from the original image. For example, the 512*512 image was obtained by deleting every other row and column from the 1024*1024 image. The 256*256 image was generated by deleting every other row and column in the 512*512 image, and so on. The number of allowed gray levels was kept at 256. These images show the dimensional proportions between various sampling densities, but their size differences make it difficult to see the effects resulting from a reduction in the number of samples. The simplest way to compare these effects is to bring all the subsampled images up to size 1024*1024 by row and column pixel replication. The results are shown in Figs. 7.2 (b) through (f). Figure7.2 (a) is the same 1024*1024, 256-level image shown in Fig.7.1; it is repeated to facilitate comparisons.

a b c
d e f
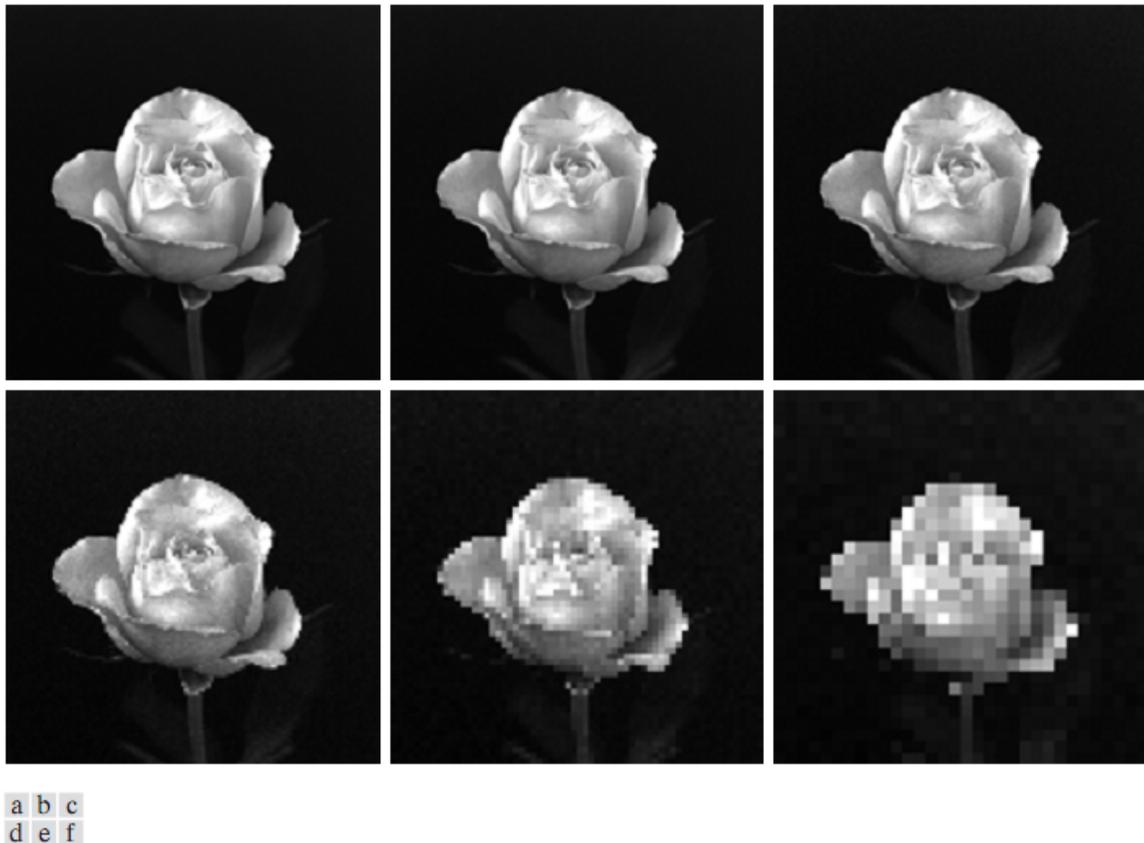
Fig. 7.2 (a) 1024*1024, 8-bit image (b) 512*512 image resampled into 1024*1024 pixels by row and column duplication (c) through (f) 256*256, 128*128, 64*64, and 32*32 images resampled into 1024*1024 pixels

Compare Fig. 7.2(a) with the 512*512 image in Fig. 7.2(b) and note that it is virtually impossible to tell these two images apart. The level of detail lost is simply too fine to be seen on the printed page at the scale in which these images are shown. Next, the 256*256 image in Fig. 7.2(c) shows a very slight fine checkerboard pattern in the borders between flower petals and the black background. A slightly more pronounced graininess throughout the image also is beginning to appear. These effects are much more visible in the 128*128 image in Fig. 7.2(d), and they become pronounced in the 64*64 and 32*32 images in Figs. 7.2 (e) and (f), respectively.

In the next example, we keep the number of samples constant and reduce the number of gray levels from 256 to 2, in integer powers of 2.Figure 7.3(a) is a 452*374 CAT projection image, displayed with k=8 (256 gray levels). Images such as this are obtained by fixing the X-ray source in one position, thus producing a 2-D image in any desired direction. Projection images are used as guides to set up the parameters for a CAT scanner, including tilt, number of slices, and range. Figures 7.3(b) through (h) were obtained by reducing the number of bits from k=7 to k=1 while keeping the spatial resolution constant at 452*374 pixels. The 256-, 128-, and 64-level images are visually identical for all practical purposes. The 32-level image shown in Fig. 7.3 (d), however, has an almost imperceptible set of very fine ridge like structures in areas of smooth gray levels (particularly in the skull).This effect, caused by the use of an insufficient number of gray levels in smooth areas of a digital image, is called false contouring, so called because the ridges resemble topographic contours in a map. False contouring generally is quite visible in images displayed using 16 or less uniformly spaced gray levels, as the images in Figs. 7.3(e) through (h) show.
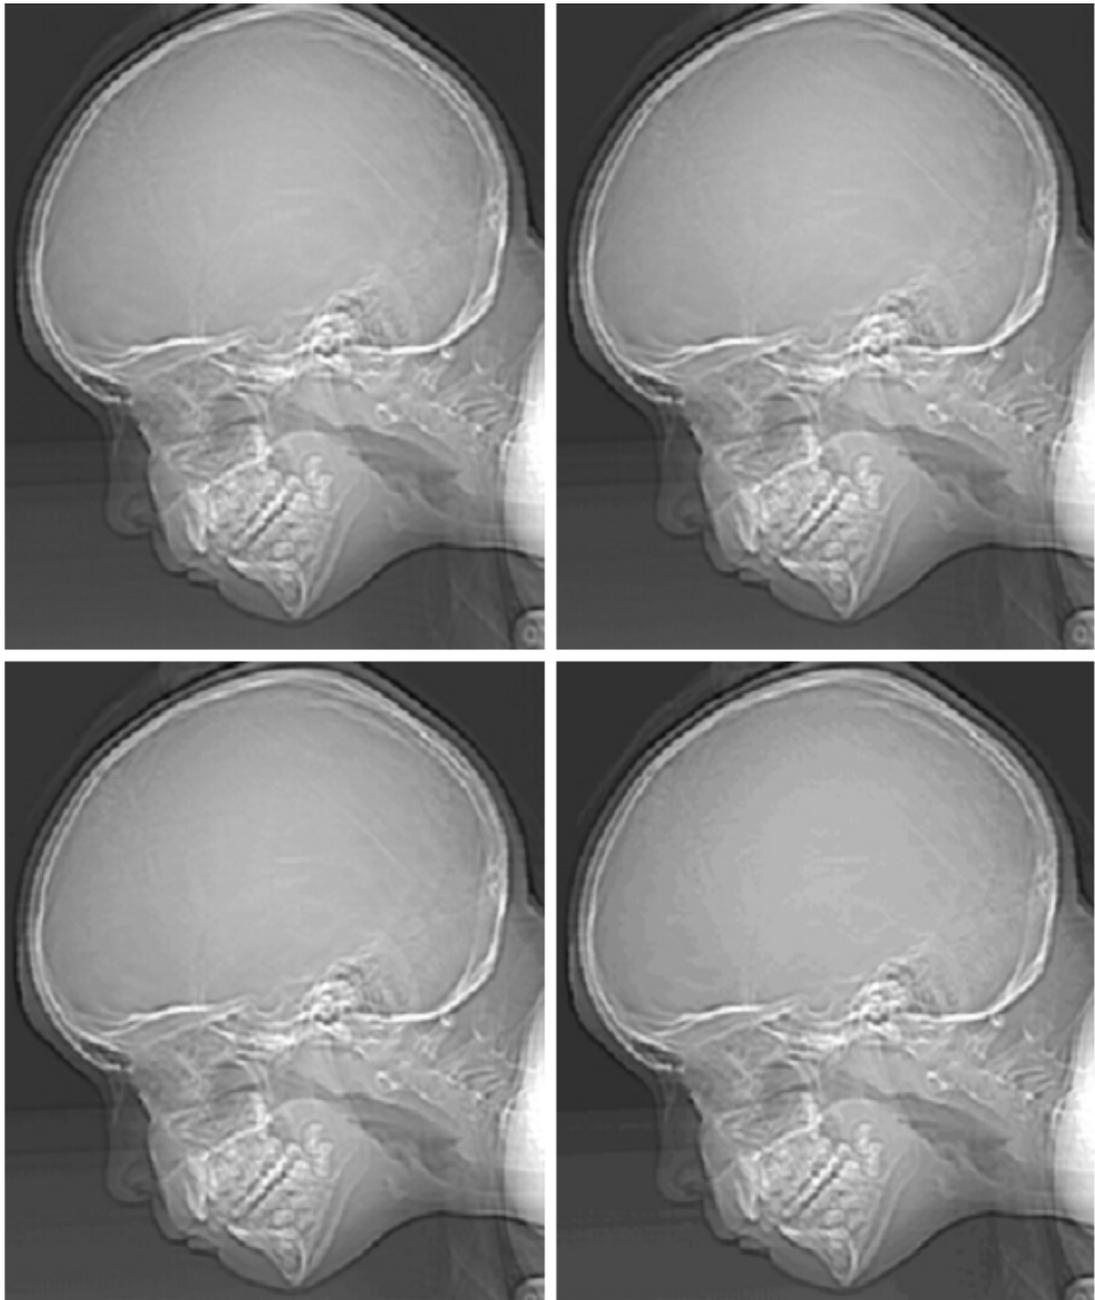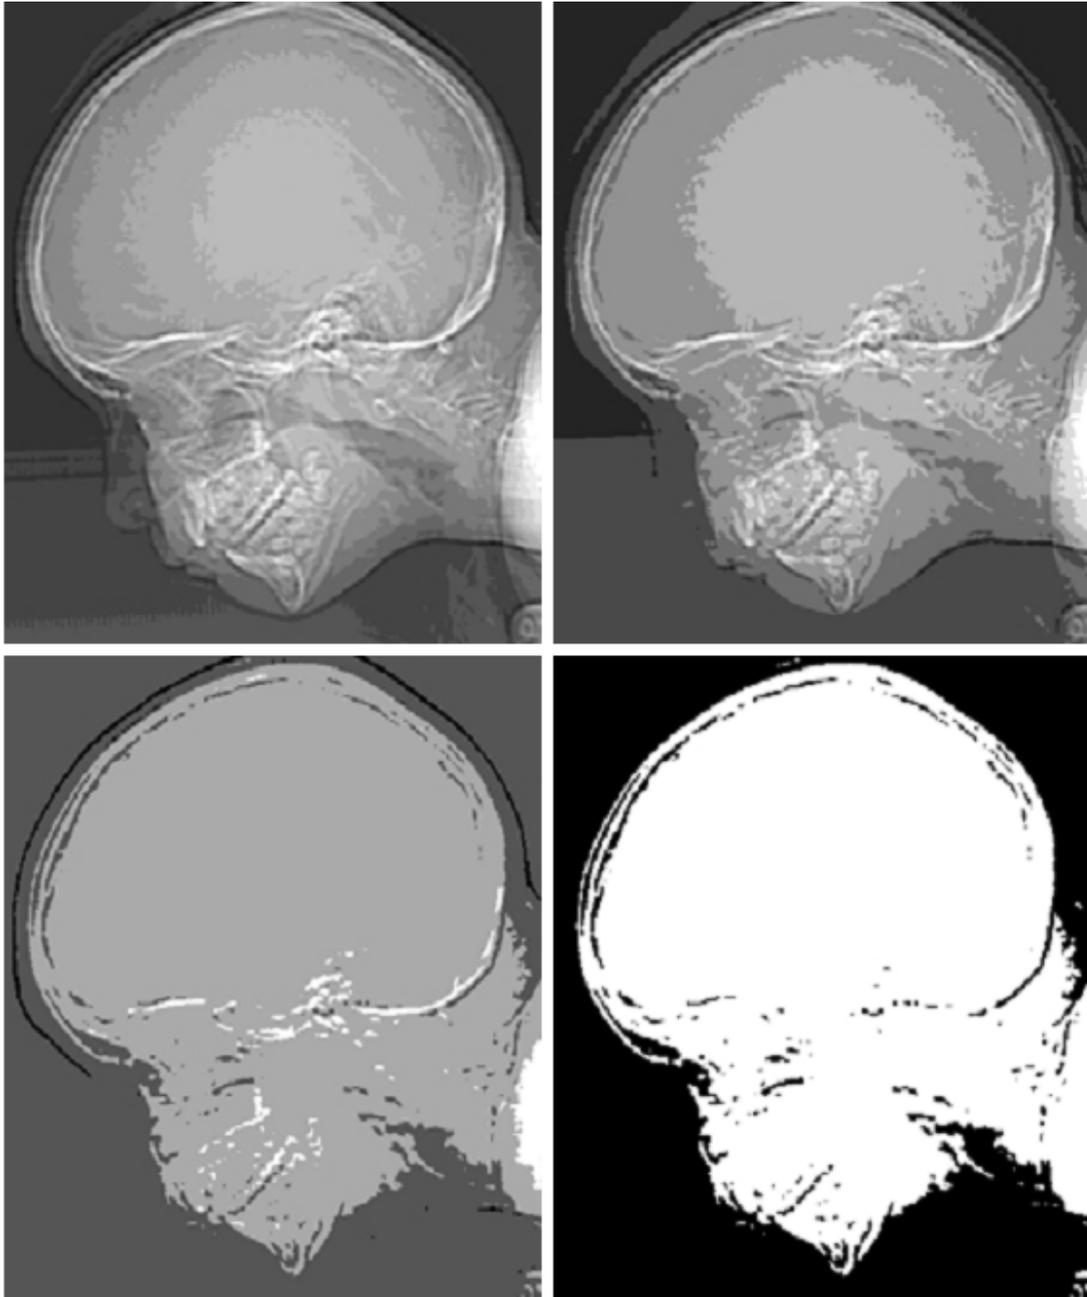
Fig. 7.3 (a) 452*374, 256-level image (b)–(d) Image displayed in 128, 64, and 32 gray levels, while keeping the spatial resolution constant (e)–(g) Image displayed in 16, 8, 4, and 2 gray levels.

As a very rough rule of thumb, and assuming powers of 2 for convenience, images of size 256*256 pixels and 64 gray levels are about the smallest images that can be expected to be reasonably free of objectionable sampling checker-boards and false contouring.

The results in Examples 7.2 and 7.3 illustrate the effects produced on image quality by varying N and k independently. However, these results only partially answer the question of how varying N and k affect images because we have not considered yet any relationships that might exist between these two parameters.

An early study by Huang [1965] attempted to quantify experimentally the effects on image quality produced by varying N and k simultaneously. The experiment consisted of a set of subjective tests. Images similar to those shown in Fig.7.4 were used. The woman's face is representative of an image with relatively little detail; the picture of the cameraman contains an intermediate amount of detail; and the crowd picture contains, by comparison, a large amount of detail. Sets of these three types of images were generated by varying N and k, and observers were then asked to rank them according to their subjective quality. Results were summarized in the form of so-called isopreference curves in the Nk-plane (Fig.7.5 shows average isopreference curves representative of curves corresponding to the images shown in Fig. 7.4).Each point in the Nk-plane represents an image having values of N and k equal to the coordinates of that point.

a b c

Fig.7.4 (a) Image with a low level of detail (b) Image with a medium level of detail (c) Image with a relatively large amount of detail

Points lying on an isopreference curve correspond to images of equal subjective quality. It was found in the course of the experiments that the isopreference curves tended to shift right and upward, but their shapes in each of the three image categories were similar to those shown in

Fig. 7.5. This is not unexpected, since a shift up and right in the curves simply means larger values for N and k, which implies better picture quality.
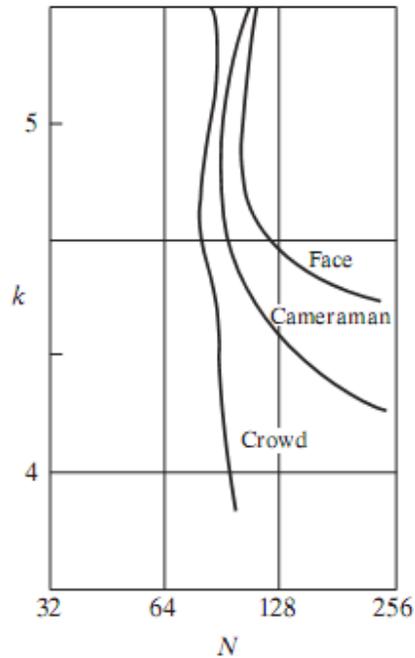
Fig.7.5. Representative isopreference curves for the three types of images in Fig.7.4

The key point of interest in the context of the present discussion is that isopreference curves tend to become more vertical as the detail in the image increases. This result suggests that for images with a large amount of detail only a few gray levels may be needed. For example, the isopreference curve in Fig.7.5 corresponding to the crowd is nearly vertical. This indicates that, for a fixed value of N, the perceived quality for this type of image is nearly independent of the number of gray levels used. It is also of interest to note that perceived quality in the other two image categories remained the same in some intervals in which the spatial resolution was increased, but the number of gray levels actually decreased. The most likely reason for this result is that a decrease in k tends to increase the apparent contrast of an image, a visual effect that humans often perceive as improved quality in an image.

8. Explain about Aliasing and Moire patterns.

Aliasing and Moiré Patterns:

Functions whose area under the curve is finite can be represented in terms of sines and cosines of various frequencies. The sine/cosine component with the highest frequency determines the highest "frequency content" of the function. Suppose that this highest frequency is finite and that the function is of unlimited duration (these functions are called band-limited functions).Then, the Shannon sampling theorem [Brace well (1995)] tells us that, if the function is sampled at a rate equal to or greater than twice its highest frequency, it is possible to recover completely the original function from its samples. If the function is undersampled, then a phenomenon called aliasing corrupts the sampled image. The corruption is in the form of additional frequency components being introduced into the sampled function. These are called aliased frequencies. Note that the sampling rate in images is the number of samples taken (in both spatial directions) per unit distance.

As it turns out, except for a special case discussed in the following paragraph, it is impossible to satisfy the sampling theorem in practice. We can only work with sampled data that are finite in duration. We can model the process of converting a function of unlimited duration into a function of finite duration simply by multiplying the unlimited function by a "gating function" that is valued 1 for some interval and 0 elsewhere. Unfortunately, this function itself has frequency components that extend to infinity. Thus, the very act of limiting the duration of a band-limited function causes it to cease being band limited, which causes it to violate the key condition of the sampling theorem. The principal approach for reducing the aliasing effects on an image is to reduce its high-frequency components by blurring the image prior to sampling. However, aliasing is always present in a sampled image. The effect of aliased frequencies can be seen under the right conditions in the form of so called Moiré patterns.

There is one special case of significant importance in which a function of infinite duration can be sampled over a finite interval without violating the sampling theorem. When a function is periodic, it may be sampled at a rate equal to or exceeding twice its highest frequency and it is possible to recover the function from its samples provided that the sampling captures exactly an integer number of periods of the function. This special case allows us to illustrate vividly the Moiré effect. Figure 8 shows two identical periodic patterns of equally spaced vertical bars, rotated in opposite directions and then superimposed on each other by multiplying the two images. A Moiré pattern, caused by a breakup of the periodicity, is seen in Fig.8 as a 2-D sinusoidal (aliased) waveform (which looks like a corrugated tin roof) running in a vertical direction. A similar pattern can appear when images are digitized (e.g., scanned) from a printed page, which consists of periodic ink dots.
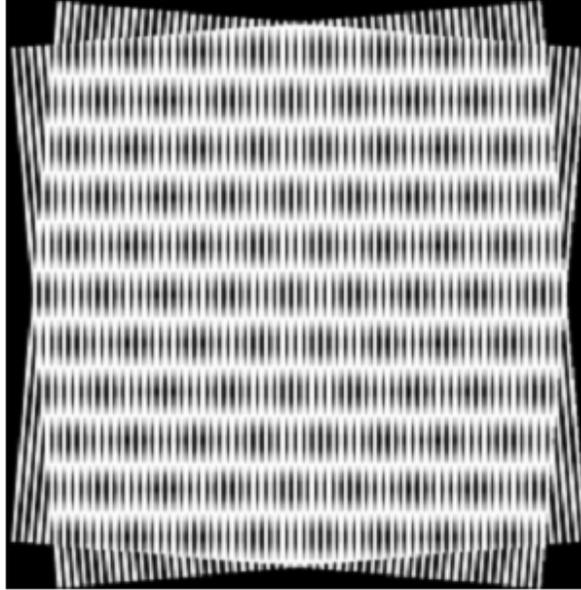
Fig.8. Illustration of the Moiré pattern effect

9. Explain about the basic relationships and distance measures between pixels in a digital image.

Neighbors of a Pixel:

A pixel p at coordinates (x, y) has four horizontal and vertical neighbors whose coordinates are given by (x+1, y), (x-1, y), (x, y+1), (x, y-1). This set of pixels, called the 4-neighbors of p, is denoted by N_4 (p). Each pixel is a unit distance from (x, y), and some of the neighbors of p lie outside the digital image if (x, y) is on the border of the image.

The four diagonal neighbors of p have coordinates (x+1, y+1), (x+1, y-1), (x-1, y+1), (x-1, y-1) and are denoted by N_D (p). These points, together with the 4-neighbors, are called the 8-neighbors of p, denoted by N_8 (p). As before, some of the points in N_D (p) and N_8 (p) fall outside the image if (x, y) is on the border of the image.

Connectivity:

Connectivity between pixels is a fundamental concept that simplifies the definition of numerous digital image concepts, such as regions and boundaries. To establish if two pixels are connected, it must be determined if they are neighbors and if their gray levels satisfy a specified criterion of similarity (say, if their gray levels are equal). For instance, in a binary image with values 0 and 1, two pixels may be 4-neighbors, but they are said to be connected only if they have the same value.

Let V be the set of gray-level values used to define adjacency. In a binary image, V={1} if we are referring to adjacency of pixels with value 1. In a grayscale image, the idea is the same, but set V typically contains more elements. For example, in the adjacency of pixels with a range of possible gray-level values 0 to 255, set V could be any subset of these 256 values. We consider three types of adjacency:

(a) 4-adjacency. Two pixels p and q with values from V are 4-adjacent if q is in the set $N_4(p)$.

(b) 8-adjacency. Two pixels p and q with values from V are 8-adjacent if q is in the set $N_8(p)$.

(c) m-adjacency (mixed adjacency).Two pixels p and q with values from V are m-adjacent if

(i) q is in $N_4(p)$, or

(ii) q is in $N_D(p)$ and the set has no pixels whose values are from V.

Mixed adjacency is a modification of 8-adjacency. It is introduced to eliminate the ambiguities that often arise when 8-adjacency is used. For example, consider the pixel arrangement shown in Fig.9 (a) for V= {1}.The three pixels at the top of Fig.9 (b) show multiple (ambiguous) 8-adjacency, as indicated by the dashed lines. This ambiguity is removed by using m-adjacency, as shown in Fig. 9 (c).Two image subsets S1 and S2 are adjacent if some pixel in S1 is adjacent to some pixel in S2. It is understood here and in the following definitions that adjacent means 4-, 8-, or m-adjacent. A (digital) path (or curve) from pixel p with coordinates (x, y) to pixel q with coordinates (s, t) is a sequence of distinct pixels with coordinates

$$(x_0, y_0), (x_1, y_1), \ldots, (x_n, y_n)$$

where $(x_0, y_0) = (x, y)$, $(x_n, y_n) = (s, t)$, and pixels (x_i, y_i) and (x_{i-1}, y_{i-1}) are adjacent for $1 \le i \le n$. In this case, n is the length of the path. If $(x_0, y_0) = (x_n, y_n)$, the path is a closed path. We can define 4-, 8-, or m-paths depending on the type of adjacency specified. For example, the paths shown in Fig. 9 (b) between the northeast and southeast points are 8-paths, and the path in Fig. 9 (c) is an m-path. Note the absence of ambiguity in the m-path. Let S represent a subset of pixels in an image. Two pixels p and q are said to be connected in S if there

exists a path between them consisting entirely of pixels in S. For any pixel p in S, the set of pixels that are connected to it in S is called a connected component of S. If it only has one connected component, then set S is called a connected set.

Let R be a subset of pixels in an image. We call R a region of the image if R is a connected set. The boundary (also called border or contour) of a region R is the set of pixels in the region that have one or more neighbors that are not in R. If R happens to be an entire image (which we recall is a rectangular set of pixels), then its boundary is defined as the set of pixels in the first and last rows and columns of the image. This extra definition is required because an image has no neighbors beyond its border. Normally, when we refer to a region, we are referring to a subset

```
0   1   1        0   1----1        0   1----1
                     |  /                |  :
0   1   0        0   1   0          0   1   0
                        `.                  `.
0   0   1        0   0   `1         0   0   `1
```
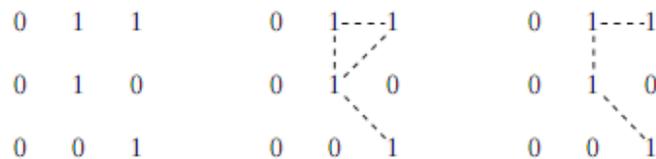
a b c

Fig.9 (a) Arrangement of pixels; (b) pixels that are 8-adjacent (shown dashed) to the center pixel; (c) m-adjacency

of an image, and any pixels in the boundary of the region that happen to coincide with the border of the image are included implicitly as part of the region boundary.

Distance Measures:

For pixels p, q, and z, with coordinates (x, y), (s, t), and (v, w), respectively, D is a distance function or metric if

(a) $D(p, q) \geq 0 \quad (D(p, q) = 0 \quad \text{iff} \quad p = q)$,
(b) $D(p, q) = D(q, p)$, and
(c) $D(p, z) \leq D(p, q) + D(q, z)$.

The **Euclidean distance** between p and q is defined as

$$D_e(p, q) = \left[(x - s)^2 + (y - t)^2\right]^{\frac{1}{2}}.$$

For this distance measure, the pixels having a distance less than or equal to some value r from(x, y) are the points contained in a disk of radius r centered at (x, y).

The **D₄ distance (also called city-block distance)** between p and q is defined as

$$D_4(p, q) = |x - s| + |y - t|.$$

In this case, the pixels having a D_4 distance from (x, y) less than or equal to some value r form a diamond centered at (x, y). For example, the pixels with D_4 distance ≤ 2 from (x, y) (the center point) form the following contours of constant distance:

```
        2
      2 1 2
    2 1 0 1 2
      2 1 2
        2
```

The pixels with D_4=1 are the 4-neighbors of (x, y).

The **D₈ distance (also called chessboard distance)** between p and q is defined as

$$D_8(p, q) = \max(|x - s|, |y - t|).$$

In this case, the pixels with D_8 distance from(x, y) less than or equal to some value r form a square centered at (x, y). For example, the pixels with D_8 distance ≤ 2 from(x, y) (the center point) form the following contours of constant 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
```

The pixels with $D_8=1$ are the 8-neighbors of (x, y). Note that the D4 and D8 distances between p and q are independent of any paths that might exist between the points because these distances involve only the coordinates of the points. If we elect to consider m-adjacency, however, the D_m distance between two points is defined as the shortest m-path between the points. In this case, the distance between two pixels will depend on the values of the pixels along the path, as well as the values of their neighbors. For instance, consider the following arrangement of pixels and assume that p, p_2 , and p_4 have value 1 and that p_1 and p_3 can have a value of 0 or 1:

$$p_3 \quad p_4$$
$$p_1 \quad p_2$$
$$p$$

Suppose that we consider adjacency of pixels valued 1 (i.e. = {1}). If p_1 and p_3 are 0, the length of the shortest m-path (the Dm distance) between p and p_4 is 2. If p1 is 1, then p_2 and p will no longer be m-adjacent (see the definition of m-adjacency) and the length of the shortest m-path becomes 3 (the path goes through the points $pp_1p_2p_4$). Similar comments apply if p_3 is 1 (and p_1 is 0); in this case, the length of the shortest m-path also is 3. Finally, if both p_1 and p_3 are 1 the length of the shortest m-path between p and p4 is 4. In this case, the path goes through the sequence of points $pp_1p_2p_3p_4$.

10. Write about perspective image transformation.

A perspective transformation (also called an imaging transformation) projects 3D points onto a plane. Perspective transformations play a central role in image processing because they provide an approximation to the manner in which an image is formed by viewing a 3D world. These transformations are fundamentally different, because they are nonlinear in that they involve division by coordinate values.

Figure 10 shows a model of the image formation process. The camera coordinate system (x, y, z) has the image plane coincident with the xy plane and the optical axis (established by the center of the lens) along the z axis. Thus the center of the image plane is at the origin, and the centre of the lens is at coordinates (0.0, λ). If the camera is in focus for distant objects, λ is the focal length of the lens. Here the assumption is that the camera coordinate system is aligned with the world coordinate system (X, Y, Z).

Let (X, Y, Z) be the world coordinates of any point in a 3-D scene, as shown in the Fig. 10. We assume throughout the following discussion that Z> λ; that is all points of interest lie in front of the lens. The first step is to obtain a relationship that gives the coordinates (x, y) of the projection of the point (X. Y, Z) onto the image plane. This is easily accomplished by the use of similar triangles. With reference to Fig. 10,
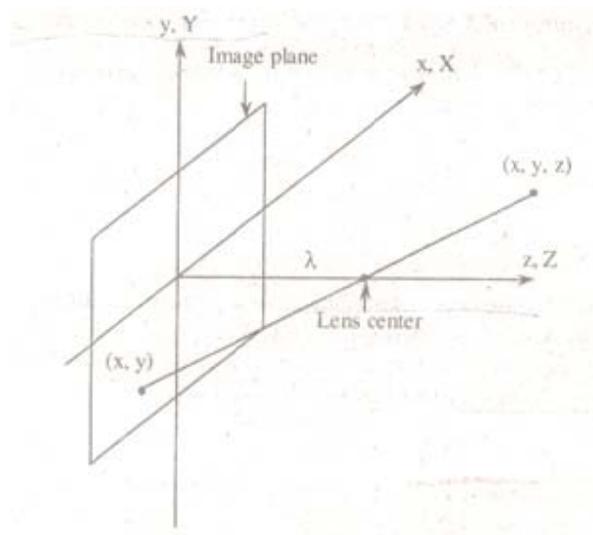
Fig.10 Basic model of the imaging process The camera coordinate system (x, y, z) is aligned with the world coordinate system (X, Y, Z)

$$\frac{x}{\lambda} = -\frac{X}{Z-\lambda}$$

$$= \frac{X}{\lambda - Z}$$

$$\frac{y}{\lambda} = -\frac{Y}{Z-\lambda}$$

$$= \frac{Y}{\lambda - Z}$$

Where the negative signs in front of X and Y indicate that image points are actually inverted, as the geometry of Fig.10 shows.
The image-plane coordinates of the projected 3-D point follow directly from above equations

$$x = \frac{\lambda X}{\lambda - Z}$$

$$y = \frac{\lambda Y}{\lambda - Z}$$

These equations are nonlinear because they involve division by the variable Z. Although we could use them directly as shown, it is often convenient to express them in linear matrix form, for rotation, translation and scaling. This is easily accomplished by using homogeneous coordinates.

The homogeneous coordinates of a point with Cartesian coordinates (X, Y, Z) are defined as (kX, kY, kZ. k), where k is an arbitrary, nonzero constant. Clearly, conversion of homogeneous coordinates back to Cartesian coordinates is accomplished by dividing the first three homogeneous coordinates by the fourth. A point in the Cartesian world coordinate system may be expressed in vector form as

$$\mathbf{w} = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

and its homogeneous counterpart is

$$\mathbf{w_h} = \begin{bmatrix} kX \\ kY \\ kZ \\ k \end{bmatrix}$$

If we define the perspective transformation matrix as

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{-1}{\lambda} & 1 \end{bmatrix}$$

The product Pw_h yields a vector denoted c_h

$$c_h = Pw_h$$

$$= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{-1}{\lambda} & 1 \end{bmatrix} \begin{bmatrix} kX \\ kY \\ kZ \\ k \end{bmatrix}$$

$$= \begin{bmatrix} kX \\ kY \\ kZ \\ \frac{-kZ}{\lambda} + k \end{bmatrix}$$

The element of c_h is the camera coordinates in homogeneous form. As indicated, these coordinates can be converted to Cartesian form by dividing each of the first three components of c_h by the fourth. Thus the Cartesian of any point in the camera coordinate system are given in vector form by

$$c = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \dfrac{\lambda X}{\lambda - Z} \\ \dfrac{\lambda Y}{\lambda - Z} \\ \dfrac{\lambda Z}{\lambda - Z} \end{bmatrix}$$

The first two components of c are the (x, y) coordinates in the image plane of a projected 3-D point (X, Y, Z). The third component is of no interest in terms of the model in Fig. 10. As shown next, this component acts as a free variable in the inverse perspective transformation

The inverse perspective transformation maps an image point back into 3-D.

$$w_h = P^{-1}c_h$$

Where P^{-1} is

$$
\mathbf{P}^{-1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{1}{\lambda} & 1 \end{bmatrix}
$$

Suppose that an image point has coordinates $(x_o, y_o, 0)$, where the 0 in the z location simply indicates that the image plane is located at $z = 0$. This point may be expressed in homogeneous vector form as

$$
\mathbf{c_h} = \begin{bmatrix} kx_o \\ ky_o \\ 0 \\ k \end{bmatrix}
$$

$$
\mathbf{w_h} = \begin{bmatrix} kx_o \\ ky_o \\ 0 \\ k \end{bmatrix}
$$

or, in Cartesian coordinates

$$
\mathbf{w} = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} x_o \\ y_o \\ 0 \end{bmatrix}
$$

This result obviously is unexpected because it gives $Z = 0$ for any 3-D point. The problem here is caused by mapping a 3-D scene onto the image plane, which is a many-to-one transformation. The image point (x_0, y_0) corresponds to the set of collinear 3-D points that lie on the line passing through $(x_o, y_o, 0)$ and $(0, 0, \lambda)$. The equation of this line in the world coordinate system; that is,

$$
X = \frac{x_o}{\lambda}(\lambda - Z)
$$

$$
Y = \frac{y_o}{\lambda}(\lambda - Z).
$$

Equations above show that unless something is known about the 3-D point that generated an image point (for example, its Z coordinate) it is not possible to completely recover the 3-D point from its image. This observation, which certainly is not unexpected, can be used to formulate the inverse perspective transformation by using the z component of c_h as a free variable instead of 0. Thus, by letting

$$c_A = \begin{bmatrix} kx_0 \\ ky_0 \\ kz \\ k \end{bmatrix}$$

It thus follows

$$w_A = \begin{bmatrix} kx_0 \\ ky_0 \\ kz \\ \dfrac{kz}{\lambda} + k \end{bmatrix}$$

which upon conversion to Cartesian coordinate gives

$$w = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} \dfrac{\lambda x_0}{\lambda + z} \\ \dfrac{\lambda y_0}{\lambda + z} \\ \dfrac{\lambda z}{\lambda + z} \end{bmatrix}$$

In other words, treating z as a free variable yields the equations

$$X = \frac{\lambda x_0}{\lambda + z}$$

$$Y = \frac{\lambda y_0}{\lambda + z}$$

$$Z = \frac{\lambda z}{\lambda + z}$$

Solving for z in terms of Z in the last equation and substituting in the first two expressions yields

$$X = \frac{x_0}{\lambda}(\lambda - Z)$$

$$Y = \frac{y_0}{\lambda}(\lambda - Z)$$

which agrees with the observation that revering a 3-D point from its image by means of the inverse perspective transformation requires knowledge of at least one of the world coordinates of the point.

Image Processing using MATLAB

by

Dr.S. Palanivel
Professor
Department of Computer Science and Engineering
Annamalai University
Annamalainagar

What is MATLAB?

- Stands for Matrix Laboratory.

- High performance tool/language for programming.

- MATLAB has collection of toolboxes such as Image Processing Toolbox, Neural Network Toolbox, Fuzzy Logic Toolbox,Signal Processing Toolbox.

- Image Processing Toolbox is a collection of MATLAB functions (Called M-Function or M-Files) used for solving real-world problems in digital image processing.

- Useful for doing Projects.

MATLAB Desktop

MATLAB Basics

- Editor/Debugger:
  >> edit filename

- For getting help
  >> doc fileame

       or

  from the symbol ? on the desktop toolbar.

- Saving and retrieving work session data
  Save workspace and Import data in the desktop toolbar.

MATLAB Function (M-File)

- M-file: Series of Matlab commands in a file
    (default extension of the file name is .m)
- >> edit name of the function

    function [outputs] = name of the function (inputs)
            ………………..
            (Matlab commands)
            ………………..

  Example:
        function [s]= negative(y)

- >> z=negative(x)

Basic Image Processing in MATLAB

- Adding path

   >>addpath D:\SDP2009\SDPimage

- Reading gray level and color images

   >> x=imread('photo1.pgm');

   >> y=imread('photo2.ppm');

- List the current variables

   >> whos

Name	Size	Bytes	Class
x	240x320	76800	uint8 array
y	240x320x3	230400	uint8 array

- Displaying the image

   >> imshow(x)

   >> figure; imshow(y)

- Writing an image

   >> imwrite(x,'D:\copyofphoto1.pgm')

Basic Image Processing in MATLAB

- Color image (RGB) to gray level image conversion

  >> z=rgb2gray(y);

- Command level help

  >> help rgb2gray

- Histogram of an image

  >> figure; imhist(z)

- Information about the image

  >> imfinfo('photo2.ppm')

- Getting pixel value of an image

  Gray value of a pixel: >> x(1,1)

  Red, green and blue values of a pixel: >> y(1,1,:)

                          or   >> y(1,1,1), y(1,1,2), y(1,1,3)

- Clear variables from memory

  >> clear

For RGB to Gray (Z = 0.299 R + 0.587 G + 0.114 B)

Basic Image Processing in MATLAB

- Finding edge of an image

    >> x=imread('photo1.pgm');

    >> y=edge(x,'canny');

    >> imshow(x);

    >> figure;imshow(y);

Basic Image Processing in MATLAB

- Finding negative of a gray level image

  ```
  >> x=imread('photo1.pgm');
  >> for i=1:240
  >> for j=1:320
  >> z(i,j)=255-double(x(i,j));
  >> end
  >> end
  >> imshow(x)
  >> figure; imshow(mat2gray(z))
  ```

Thank You

1.What is meant by image enhancement by point processing? Discuss any two methods in it.

Basic Gray Level Transformations:

The study of image enhancement techniques is done by discussing gray-level transformation functions. These are among the simplest of all image enhancement techniques. The values of pixels, before and after processing, will be denoted by r and s, respectively. As indicated in the previous section, these values are related by an expression of the form s=T(r), where T is a transformation that maps a pixel value r into a pixel value s. Since we are dealing with digital quantities, values of the transformation function typically are stored in a one-dimensional array and the mappings from r to s are implemented via table lookups. For an 8-bit environment, a lookup table containing the values of T will have 256 entries. As an introduction to gray-level transformations, consider Fig. 1.1, which shows three basic types of functions used frequently for image enhancement: linear (negative and identity transformations), logarithmic (log and inverse-log transformations), and power-law (nth power and nth root transformations).The identity function is the trivial case in which output intensities are identical to input intensities. It is included in the graph only for completeness.

Image Negatives:

The negative of an image with gray levels in the range [0, L-1] is obtained by using the negative transformation shown in Fig.1.1, which is given by the expression

$$s = L - 1 - r.$$

Reversing the intensity levels of an image in this manner produces the equivalent of a photographic negative. This type of processing is particularly suited for enhancing white or gray detail embedded in dark regions of an image, especially when the black areas are dominant in size.
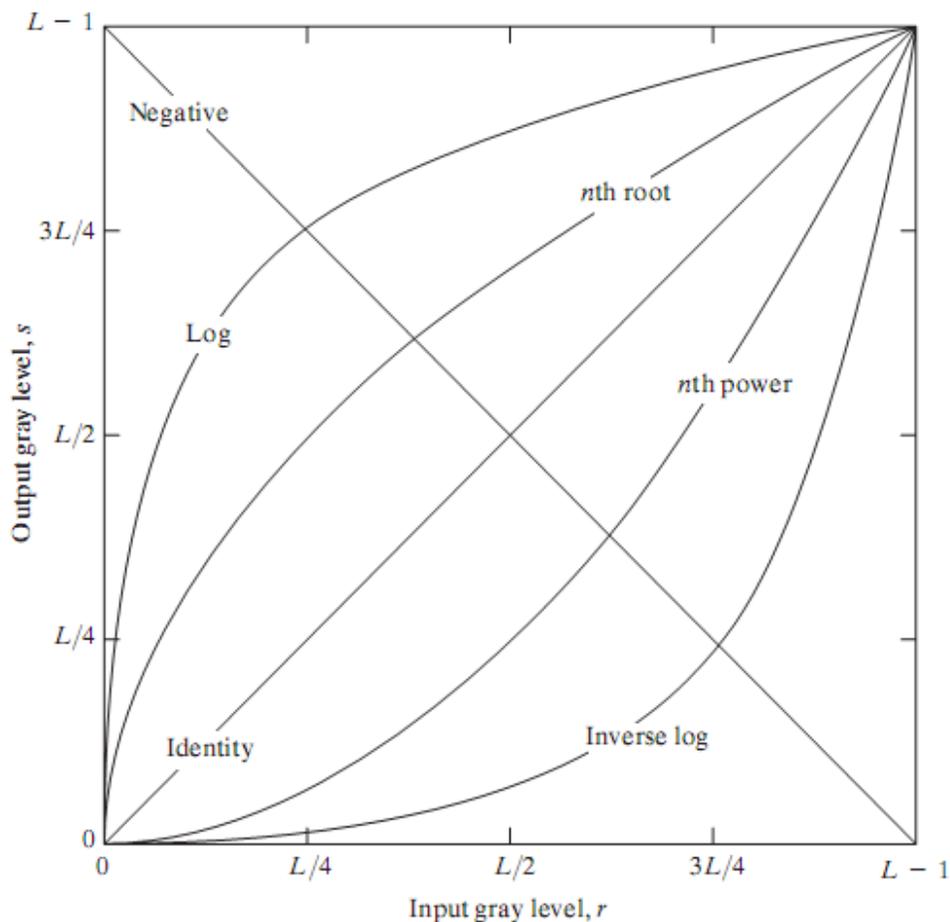
Fig.1.1 Some basic gray-level transformation functions used for image enhancement

 Log Transformations:

The general form of the log transformation shown in Fig.1.1 is

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

where c is a constant, and it is assumed that $r \geq 0$. The shape of the log curve in Fig. 1.1 shows that this transformation maps a narrow range of low gray-level values in the input image into a wider range of output levels. The opposite is true of higher values of input levels. We would use a transformation of this type to expand the values of dark pixels in an image while compressing the higher-level values. The opposite is true of the inverse log transformation.

Any curve having the general shape of the log functions shown in Fig. 1.1 would accomplish this spreading/compressing of gray levels in an image. In fact, the power-law transformations discussed in the next section are much more versatile for this purpose than the log transformation. However, the log function has the important characteristic that it compresses the dynamic range of images with large variations in pixel values. A classic illustration of an application in which pixel values have a large dynamic range is the Fourier spectrum. At the moment,we are concerned only with the image characteristics of spectra. It is not unusual to encounter spectrum values that range from 0 to or higher.While processing numbers such as these presents no problems for a computer, image display systems generally will not be able to reproduce faithfully such a wide range of intensity values. The net effect is that a significant degree of detail will be lost in the display of a typical Fourier spectrum.

Power-Law Transformations:

Power-law transformations have the basic form

$$s = cr^{\gamma}$$

where c and g are positive constants. Sometimes Eq. is written as $s = c(r + \varepsilon)^{\gamma}$

to account for an offset (that is, a measurable output when the input is zero).However, offsets typically are an issue of display calibration and as a result they are normally ignored in Eq. Plots of s versus r for various values of g are shown in Fig. 1.2. As in the case of the log transformation, power-law curves with fractional values of g map a narrow range of dark input values into a wider range of output values,with the opposite being true for high-er values of input levels. Unlike the log function, however, we notice here a family of possible transformation curves obtained simply by varying γ. As expected, we see in Fig.1.2 that curves generated with values of g>1 have exactly the opposite effect as those generated with values of g<1. Finally, we note that Eq. reduces to the identity transformation when c = γ = 1. A variety of devices used for image capture, printing, and display respond according to a power law.By convention, the exponent in the power-law equation is referred to as gamma. The proces used to correct this power-law response phenomena is called gamma correction. For example, cathode ray tube (CRT) devices have an intensity-to-voltage response that is a power function, with exponents varying from approximately 1.8 to 2.5.With reference to the curve for g=2.5 in Fig.1.2, we see that such display systems would tend to produce images that are darker than intended.
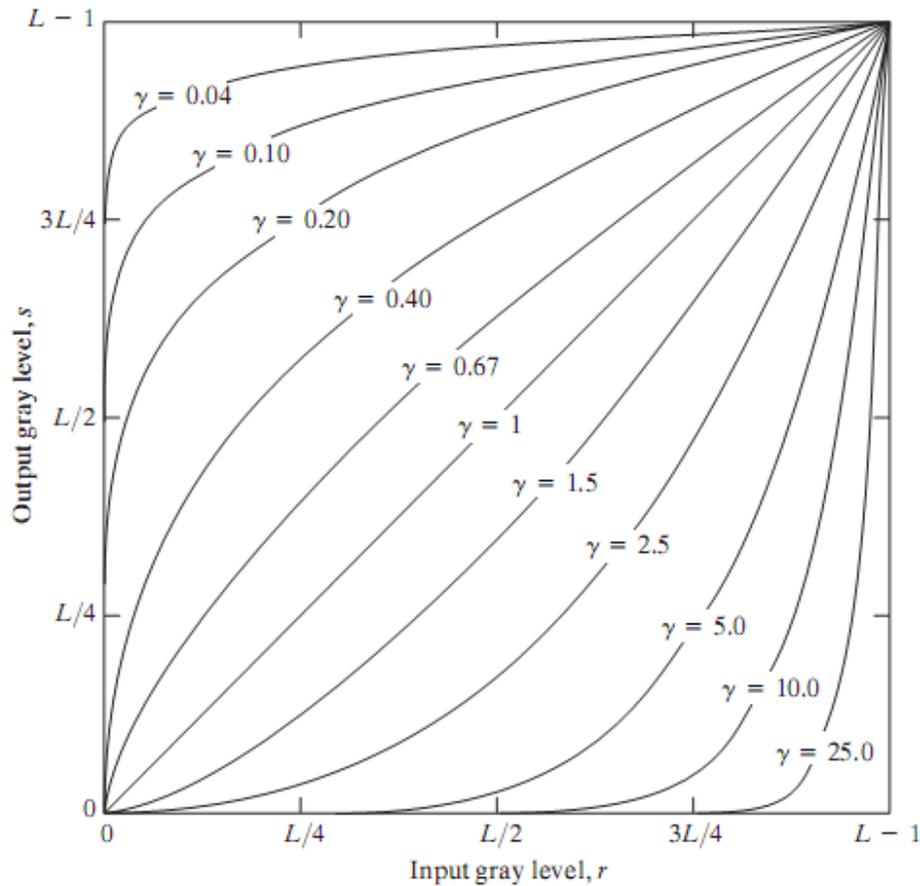
Fig.1.2 Plots of the equation $s = cr^{\gamma}$ **for various values of** γ **(c=1 in all cases).**

Piecewise-Linear Transformation Functions:

The principal advantage of piecewise linear functions over the types of functions we have discussed above is that the form of piecewise functions can be arbitrarily complex. In fact, as we will see shortly, a practical implementation of some important transformations can be formulated only as piecewise functions. The principal disadvantage of piecewise functions is that their specification requires considerably more user input.

Contrast stretching:

One of the simplest piecewise linear functions is a contrast-stretching transformation. Low-contrast images can result from poor illumination, lack of dynamic range in the imaging sensor, or even wrong setting of a lens aperture during image acquisition.The idea behind contrast stretching is to increase the dynamic range of the gray levels in the image being processed.

Figure 1.3 (a) shows a typical transformation used for contrast stretching.

The locations of points (r_1, s_1) and (r_2, s_2) control the shape of the transformation
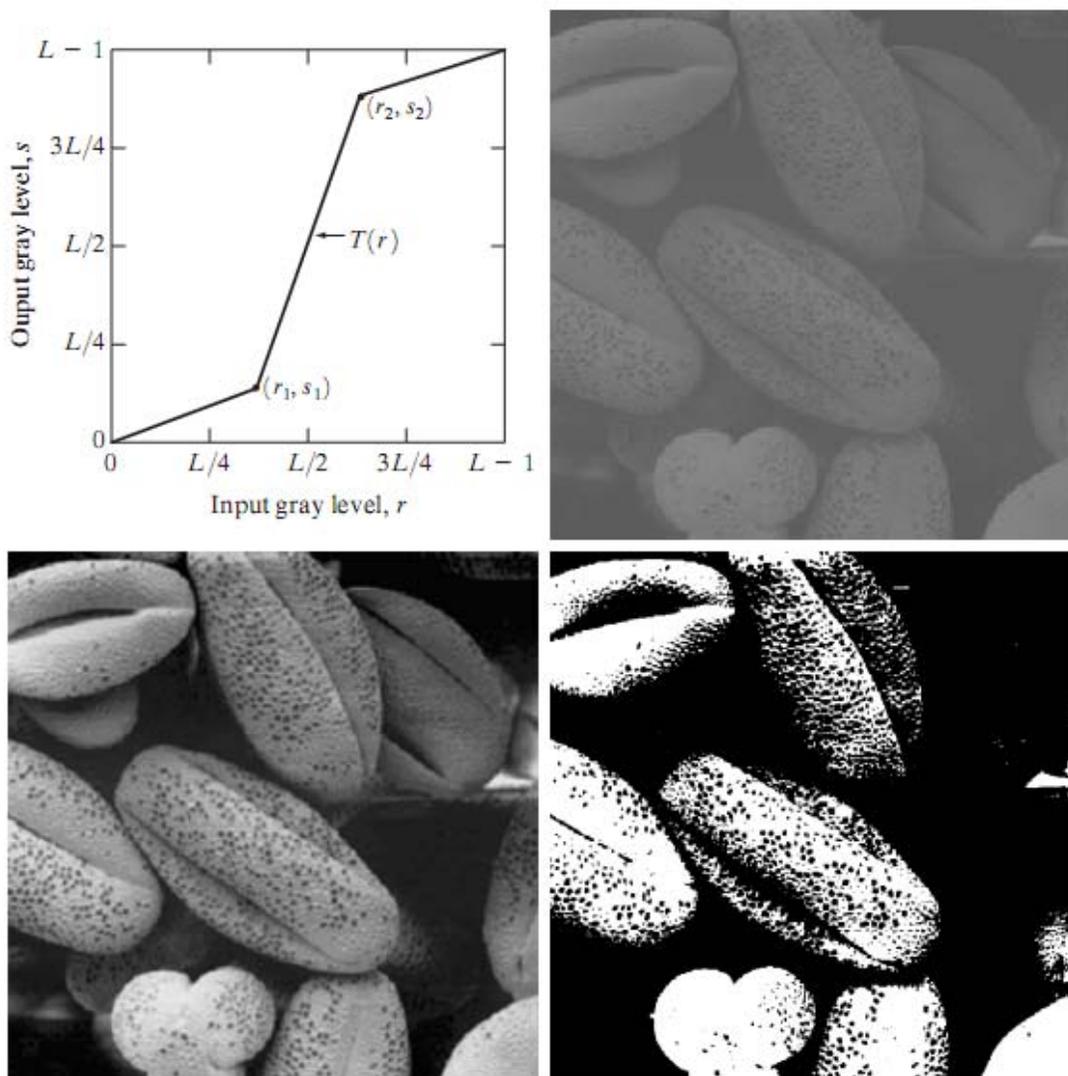
Fig.1.3 Contrast Stretching (a) Form of Transformation function (b) A low-contrast image (c) Result of contrast stretching (d) Result of thresholding.

function. If r1=s1 and r2=s2, the transformation is a linear function that produces no changes in gray levels. If r1=r2,s1=0 and s2=L-1, the transformation becomes a thresholding function that creates a binary image, as illustrated in Fig. 1.3 (b). Intermediate values of (r_1, s_1) and (r_2, s_2) produce various degrees of spread in the gray levels of the output image, thus affecting its contrast. In general, $r1 \leq r2$ and $s1 \leq s2$ is assumed so that the function is single valued and monotonically increasing.This condition preserves the order of gray levels, thus preventing the creation of intensity artifacts in the processed image.

Figure 1.3 (b) shows an 8-bit image with low contrast. Fig. 1.3(c) shows the result of contrast stretching, obtained by setting $(r_1, s_1) = (r_{min}, 0)$ and $(r_2, s_2) = (r_{max}, L-1)$ where r_{min} and r_{max} denote the minimum and maximum gray levels in the image, respectively.Thus, the transformation function stretched the levels linearly from their original range to the full range [0, L-1]. Finally, Fig. 1.3 (d) shows the result of using the thresholding function defined previously,with r1 = r2 = m, the mean gray level in the image.The original image on which these results are based is a scanning electron microscope image of pollen,magnified approximately 700 times.

Gray-level slicing:

Highlighting a specific range of gray levels in an image often is desired. Applications include enhancing features such as masses of water in satellite imagery and enhancing flaws in X-ray images.There are several ways of doing level slicing, but most of them are variations of two basic themes.One approach is to display a high value for all gray levels in the range of interest and a low value for all other gray levels.This transformation, shown in Fig. 1.4 (a), produces a binary image.The second approach, based on the transformation shown in Fig. 1.4 (b), brightens the desired range of gray levels but preserves the background and gray-level tonalities in the image. Figure 1.4(c) shows a gray-scale image, and Fig. 1.4 (d) shows the result of using the transformation in Fig. 1.4 (a).Variations of the two transformations shown in Fig. 1.4 are easy to formulate.
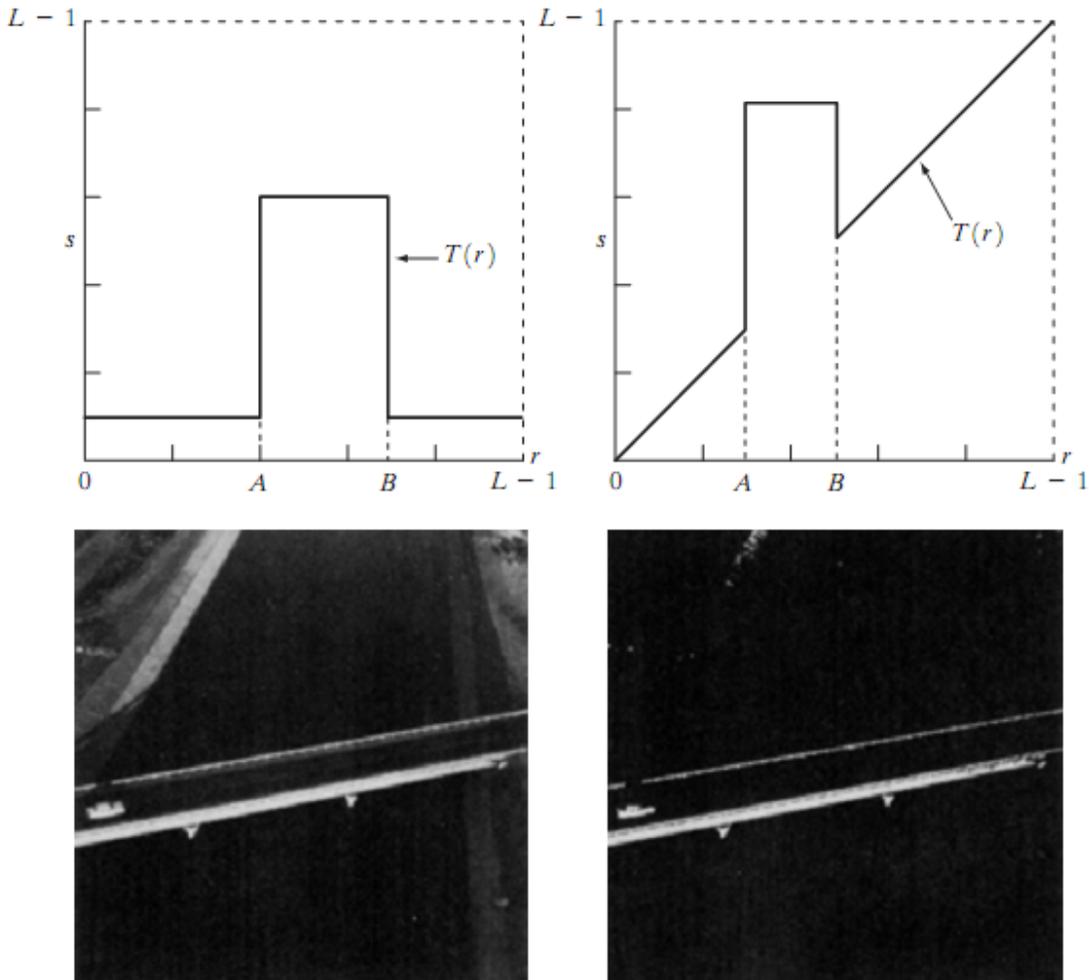
Fig.1.4 (a) This transformation highlights range [A, B] of gray levels and reduce all others to a constant level (b) This transformation highlights range [A, B] but preserves all other levels (c) An image (d) Result of using the transformation in (a).

Bit-plane slicing:

Instead of highlighting gray-level ranges, highlighting the contributionmade to total image appearance by specific bits might be desired. Suppose that each pixel in an image is represented by 8 bits. Imagine that the image is composed of eight 1-bit planes, ranging from bit-plane 0 for the least significant bit to bit plane 7 for the most significant bit. In terms of 8-bit bytes, plane 0

contains all the lowest order bits in the bytes comprising the pixels in the image and plane 7 contains all the high-order bits.Figure 1.5 illustrates these ideas, and Fig. 1.7 shows the various bit planes for the image shown in Fig.1.6 . Note that the higher-order bits (especially the top four) contain themajority of the visually significant data.The other bit planes contribute tomore subtle details in the image. Separating a digital image into its bit planes is useful for analyzing the relative importance played by each bit of the image, a process that aids in determining the adequacy of the number of bits used to quantize each pixel.
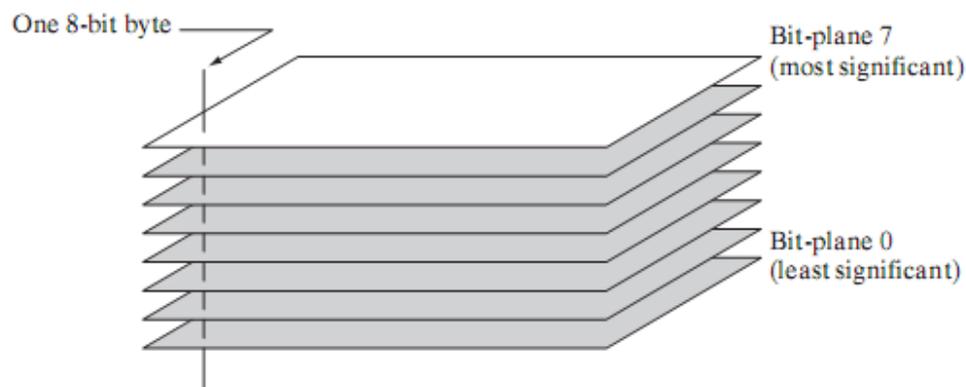
Fig.1.5 Bit-plane representation of an 8-bit image.

In terms of bit-plane extraction for an 8-bit image, it is not difficult to show that the (binary) image for bit-plane 7 can be obtained by processing the input image with a thresholding gray-level transformation function that (1) maps all levels in the image between 0 and 127 to one level (for example, 0); and (2) maps all levels between 129 and 255 to another (for example, 255).
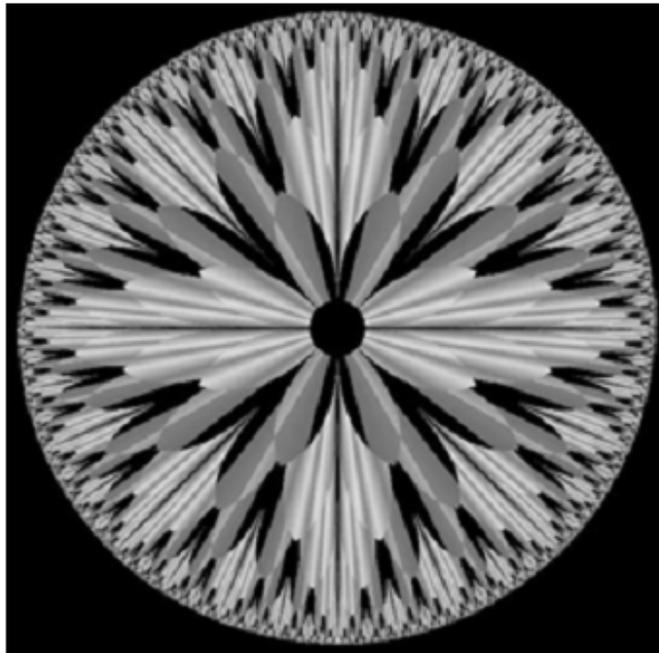
Fig.1.6 An 8-bit fractal image
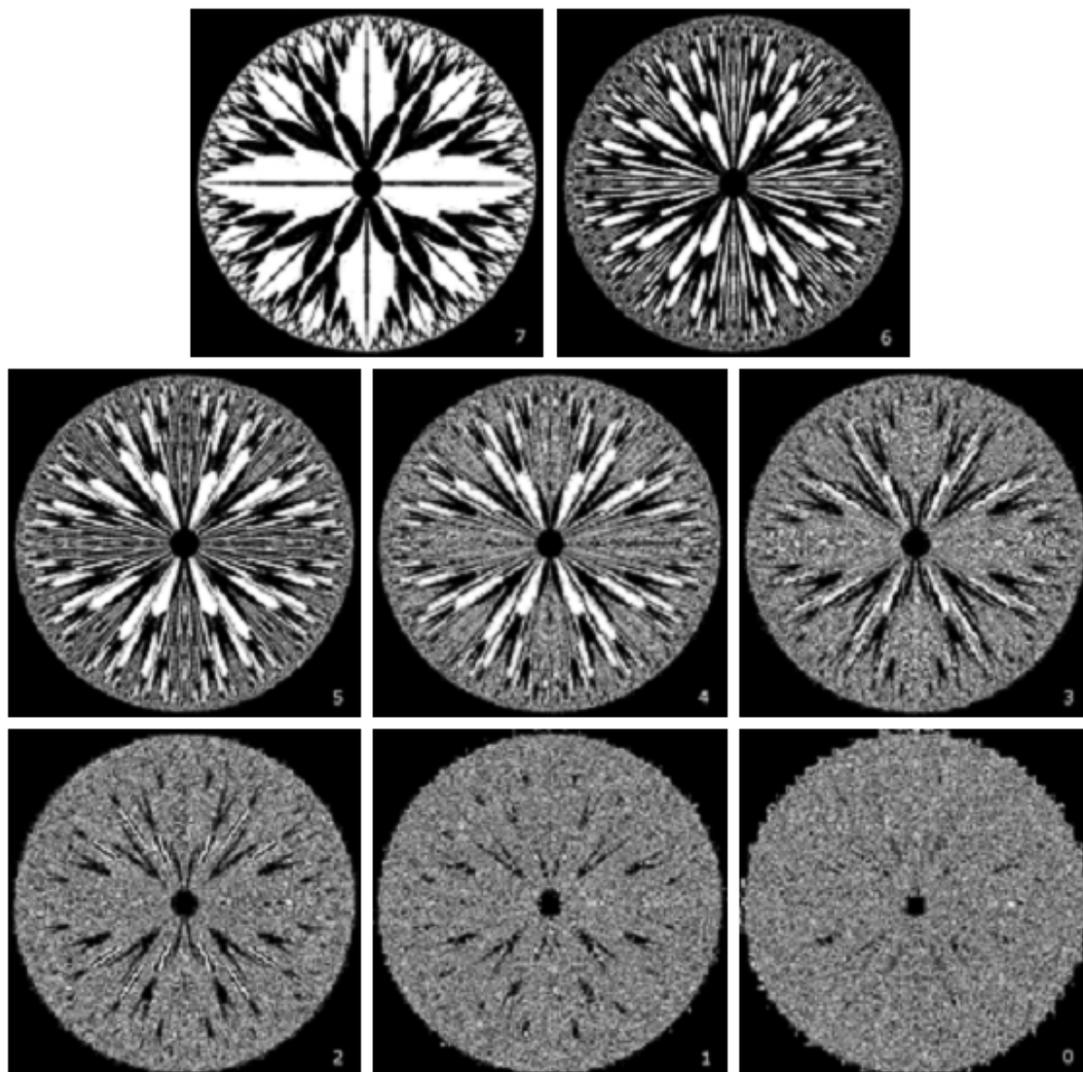
Fig.1.7 The eight bit planes of the image in Fig.1.6. The number at the bottom, right of each image identifies the bit plane.

2. What is the objective of image enhancement. Define spatial domain. Define point processing.

The term spatial domain refers to the aggregate of pixels composing an image. Spatial domain methods are procedures that operate directly on these pixels. Spatial domain processes will be denoted by the expression

$$g(x, y) = T[f(x, y)]$$

where f(x, y) is the input image, g(x, y) is the processed image, and T is an operator on f, defined over some neighborhood of (x, y). In addition,T can operate on a set of input images, such as performing the pixel-by-pixel sum of K images for noise reduction.

The principal approach in defining a neighborhood about a point (x, y) is to use a square or rectangular subimage area centered at (x, y), as Fig.2.1 shows. The center of the subimage is moved from pixel to pixel starting, say, at the top left corner. The operator T is applied at each location (x, y) to yield the output, g, at that location.The process utilizes only the pixels in the area of the image spanned by the neighborhood.
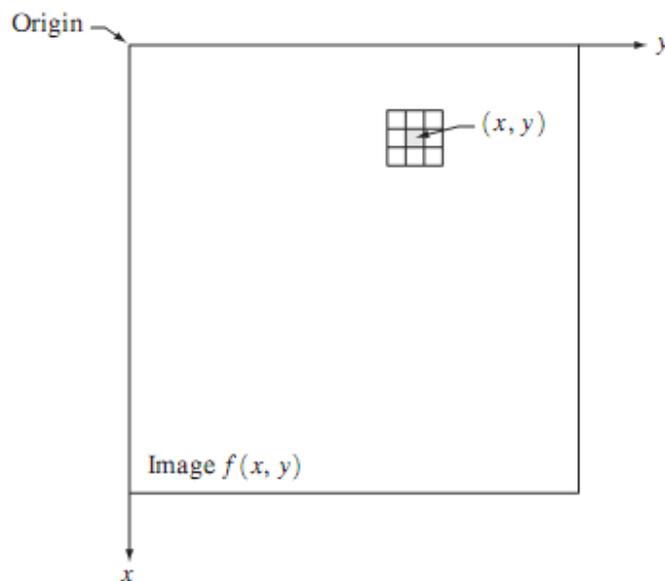
Fig.2.1 A 3*3 neighborhood about a point (x, y) in an image.

Although other neighborhood shapes, such as approximations to a circle, sometimes are used, square and rectangular arrays are by far the most predominant because of their ease of implementation. The simplest form of T is when the neighborhood is of size 1*1 (that is, a single pixel). In this case, g depends only on the value of f at (x, y), and T becomes a gray-level (also called an intensity or mapping) transformation function of the form

$$s = T(r)$$

where, for simplicity in notation, r and s are variables denoting, respectively, the gray level of f(x, y) and g(x, y) at any point (x, y). For example, if T(r) has the form shown in Fig. 2.2(a), the effect of this transformation would be to produce an image of higher contrast than the original by darkening the levels below m and brightening the levels above m in the original image. In this technique, known as contrast stretching, the values of r below m are compressed by the transformation function into a narrow range of s, toward black.The opposite effect takes place for values of r above m. In the limiting case shown in Fig. 2.2(b), T(r) produces a two-level (binary) image. A mapping of this form is called a thresholding function. Some fairly simple, yet powerful, processing approaches can be formulated with gray-level transformations. Because enhancement at any point in an image depends only on the gray level at that point, techniques in this category often are referred to as point processing.
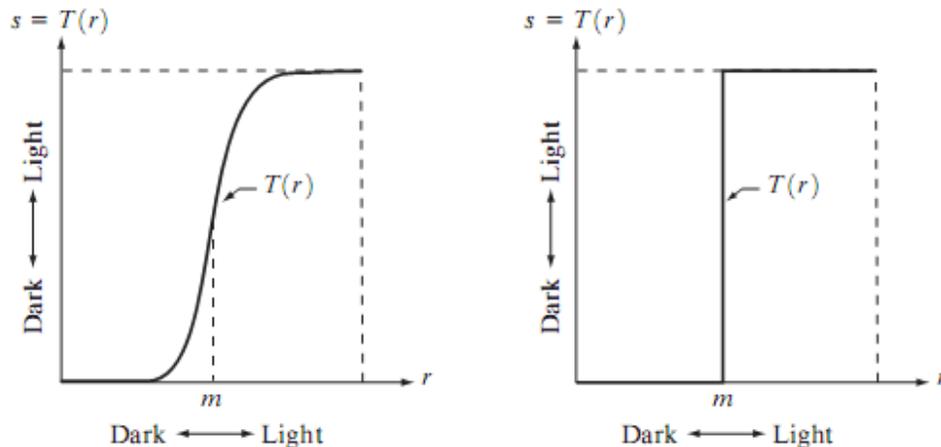
Fig.2.2 Graylevel transformation functions for contrast enhancement.

Larger neighborhoods allow considerably more flexibility. The general approach is to use a function of the values of f in a predefined neighborhood of (x, y) to determine the value of g at (x, y).One of the principal approaches in this formulation is based on the use of so-called masks

(also referred to as filters, kernels, templates, or windows). Basically, a mask is a small (say, 3*3) 2-D array, such as the one shown in Fig. 2.1, in which the values of the mask coefficients determine the nature of the process, such as image sharpening.

3. Define histogram of a digital image. Explain how histogram is useful in image enhancement?

Histogram Processing:

The histogram of a digital image with gray levels in the range [0, L-1] is a discrete function $h(r_k)$ = (n_k), where r_k is the kth gray level and n_k is the number of pixels in the image having gray level r_k. It is common practice to normalize a histogram by dividing each of its values by the total number of pixels in the image, denoted by n. Thus, a normalized histogram is given by

$$p(r_k) = n_k/n$$

for k=0,1,…… .,L-1. Loosely speaking, $p(r_k)$ gives an estimate of the probability of occurrence of gray level r_k. Note that the sum of all components of a normalized histogram is equal to 1.

Histograms are the basis for numerous spatial domain processing techniques.Histogram manipulation can be used effectively for image enhancement. Histograms are simple to calculate in software and also lend themselves to economic hardware implementations, thus making them a popular tool for real-time image processing.

As an introduction to the role of histogram processing in image enhancement, consider Fig. 3, which is the pollen image shown in four basic gray-level characteristics: dark, light, low contrast, and high contrast.The right side of the figure shows the histograms corresponding to these images. The horizontal axis of each histogram plot corresponds to gray level values, r_k.

The vertical axis corresponds to values of $h(r_k)$ = n_k or $p(r_k)$ = n_k/n if the values are normalized.Thus, as indicated previously, these histogram plots are simply plots of $h(r_k)$ = n_k versus r_k or $p(r_k)$ = n_k/n versus r_k.
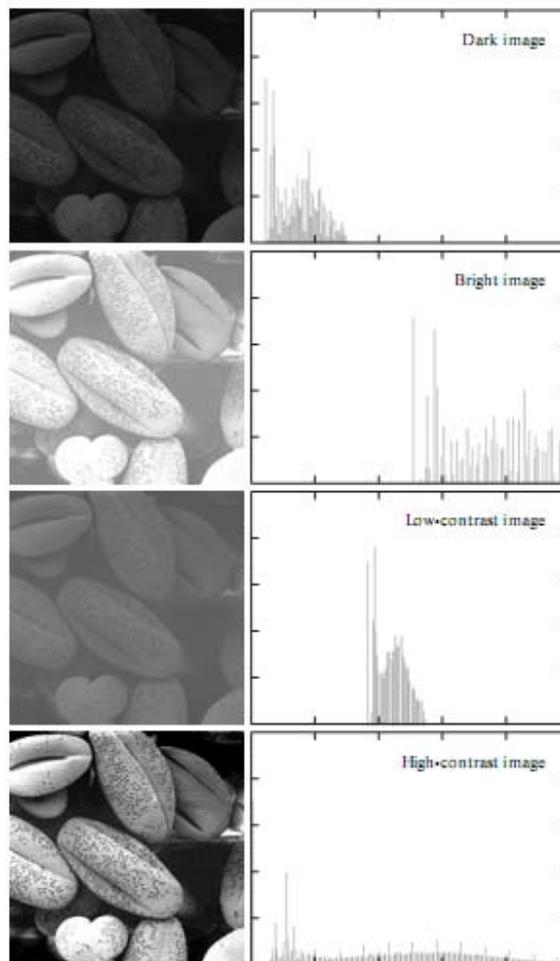
Fig.3 Four basic image types: dark, light, low contrast, high contrast, and their corresponding histograms.

We note in the dark image that the components of the histogram are concentrated on the low (dark) side of the gray scale. Similarly, the components of the histogram of the bright image are biased toward the high side of the gray scale.An image with low contrast has a histogram that will be narrow and will be centered toward the middle of the gray scale. For a monochrome image this implies a dull,washed-out gray look. Finally,we see that the components of the histogram in the high-contrast image cover a broad range of the gray scale and, further, that the distribution of pixels is not too far from uniform,with very few vertical lines being much higher than the others. Intuitively, it is reasonable to conclude that an image whose pixels tend to occupy the entire range of possible gray levels and, in addition, tend to be distributed uniformly,will have an appearance of high contrast and will exhibit a large variety of gray tones. The net effect will be an image that shows a great deal of gray-level detail and has high dynamic

range. It will be shown shortly that it is possible to develop a transformation function that can automatically achieve this effect, based only on information available in the histogram of the input image.

4. Write about histogram equalization.

Histogram Equalization:

Consider for a moment continuous functions, and let the variable r represent the gray levels of the image to be enhanced. We assume that r has been normalized to the interval [0, 1], with r=0 representing black and r=1 representing white. Later, we consider a discrete formulation and allow pixel values to be in the interval [0, L-1]. For any r satisfying the aforementioned conditions, we focus attention on transformations of the form

$$s = T(r) \qquad 0 \le r \le 1$$

that produce a level s for every pixel value r in the original image. For reasons that will become obvious shortly, we assume that the transformation function T(r) satisfies the following conditions:

(a) T(r) is single-valued and monotonically increasing in the interval $0 \le r \le 1$; and

(b) $0 \le T(r) \le 1$ for $0 \le r \le 1$.

The requirement in (a) that T(r) be single valued is needed to guarantee that the inverse transformation will exist, and the monotonicity condition preserves the increasing order from black to white in the output image. A transformation function that is not monotonically increasing could result in at least a section of the intensity range being inverted, thus producing some inverted gray levels in the output image. Finally, condition (b) guarantees that the output gray levels will be in the same range as the input levels. Figure 4.1 gives an example of a transformation function that satisfies these two conditions. The inverse transformation from s back to r is denoted

$$r = T^{-1}(s) \qquad 0 \le s \le 1.$$

It can be shown by example that even if T(r) satisfies conditions (a) and (b), it is possible that the corresponding inverse $T^{-1}(s)$ may fail to be single valued.
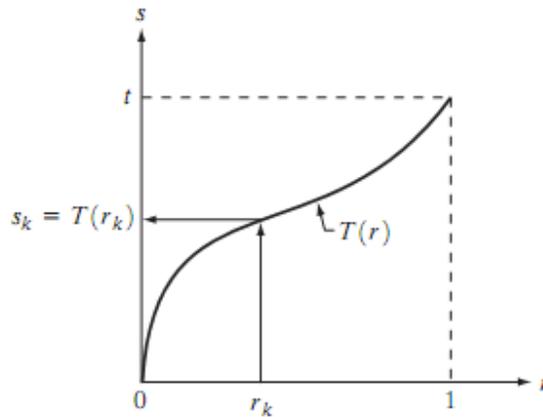
Fig.4.1 A gray-level transformation function that is both single valued and monotonically increasing.

The gray levels in an image may be viewed as random variables in the interval [0, 1].One of the most fundamental descriptors of a random variable is its probability density function (PDF).Let $p_r(r)$ and $p_s(s)$ denote the probability density functions of random variables r and s, respectively,where the subscripts on p are used to denote that p_r and p_s are different functions.A basic result from an elementary probability theory is that, if $p_r(r)$ and T(r) are known and T^{-1} (s) satisfies condition (a), then the probability density function $p_s(s)$ of the transformed variable s can be obtained using a rather simple formula:

$$p_s(s) = p_r(r)\left|\frac{dr}{ds}\right|.$$

Thus, the probability density function of the transformed variable, s, is determined by the gray-level PDF of the input image and by the chosen transformation function. A transformation function of particular importance in image processing has the form

$$s = T(r) = \int_0^r p_r(w)\, dw$$

where w is a dummy variable of integration.The right side of Eq. above is recognized as the cumulative distribution function (CDF) of random variable r. Since probability density functions are always positive, and recalling that the integral of a function is the area under the function, it follows that this transformation function is single valued and monotonically increasing, and, therefore, satisfies condition (a). Similarly, the integral of a probability density function for variables in the range [0, 1] also is in the range [0, 1], so condition (b) is satisfied as well.

Given transformation function T(r),we find $p_s(s)$ by applying Eq. We know from basic calculus (Leibniz's rule) that the derivative of a definite integral with respect to its upper limit is simply the integrand evaluated at that limit. In other words,

$$\frac{ds}{dr} = \frac{dT(r)}{dr}$$

$$= \frac{d}{dr}\left[\int_0^r p_r(w)\, dw\right]$$

$$= p_r(r).$$

Substituting this result for dr/ds, and keeping in mind that all probability values are positive, yields

$$p_s(s) = p_r(r)\left|\frac{dr}{ds}\right|$$

$$= p_r(r)\left|\frac{1}{p_r(r)}\right|$$

$$= 1 \qquad 0 \le s \le 1.$$

Because $p_s(s)$ is a probability density function, it follows that it must be zero outside the interval [0, 1] in this case because its integral over all values of s must equal 1.We recognize the form of $p_s(s)$ as a uniform probability density function. Simply stated, we have demonstrated that performing the transformation function yields a random variable s characterized by a uniform probability density function. It is important to note from Eq. discussed above that T(r) depends on $p_r(r)$, but, as indicated by Eq. after it, the resulting $p_s(s)$ always is uniform, independent of the form of $p_r(r)$. For discrete values we deal with probabilities and summations instead of probability density functions and integrals. The probability of occurrence of gray level r in an image is approximated by

$$p_r(r_k) = \frac{n_k}{n} \qquad k = 0, 1, 2, \ldots, L - 1$$

where, as noted at the beginning of this section, n is the total number of pixels in the image, n_k is the number of pixels that have gray level r_k, and L is the total number of possible gray levels in the image.The discrete version of the transformation function given in Eq. is

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

$$= \sum_{j=0}^{k} \frac{n_j}{n} \qquad k = 0, 1, 2, \ldots, L - 1.$$

Thus, a processed (output) image is obtained by mapping each pixel with level r_k in the input image into a corresponding pixel with level s_k in the output image. As indicated earlier, a plot of $p_r(r_k)$ versus r_k is called a histogram. The transformation (mapping) is called histogram equalization or histogram linearization. It is not difficult to show that the transformation in Eq. satisfies conditions (a) and (b) stated previously. Unlike its continuos counterpart, it cannot be proved in general that this discrete transformation will produce the discrete equivalent of a uniform probability density function, which would be a uniform histogram.
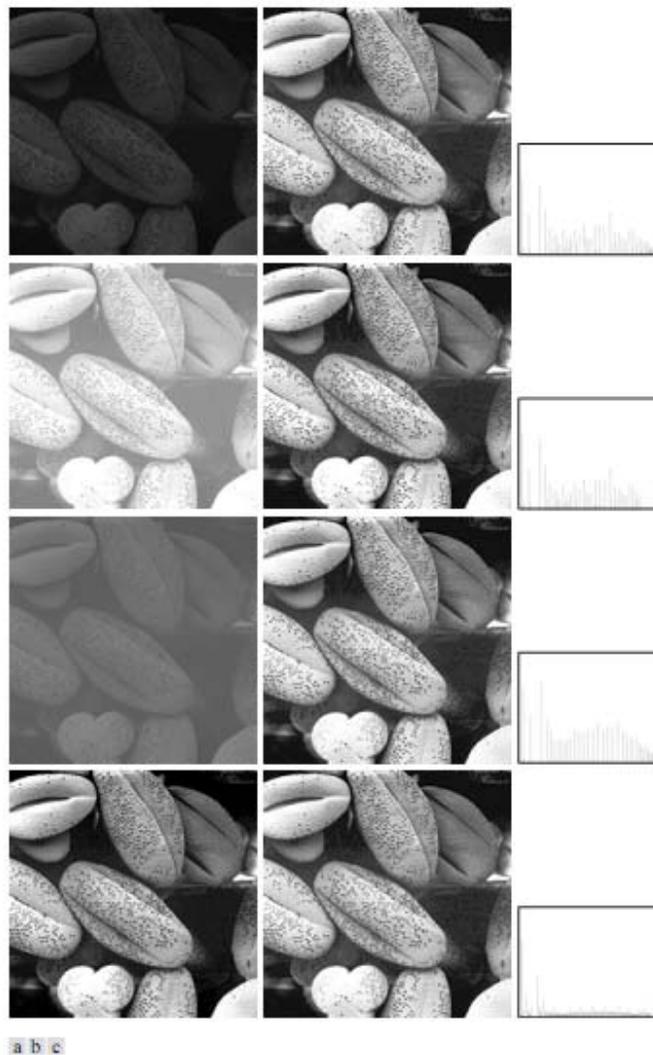
Fig.4.2   (a) Images from Fig.3 (b) Results of histogram equalization. (c) Corresponding histograms.

The inverse transformation from s back to r is denoted by

$$r_k = T^{-1}(s_k) \qquad k = 0, 1, 2, \ldots, L - 1$$

5. Write about histogram specification.

Histogram Matching (Specification):

Histogram equalization automatically determines a transformation function that seeks to produce an output image that has a uniform histogram.When automatic enhancement is desired, this is a good approach because the results from this technique are predictable and the method is simple to implement. In particular, it is useful sometimes to be able to specify the shape of the histogram that we wish the processed image to have.The method used to generate a processed image that has a specified histogram is called histogram matching or histogram specification.

Development of the method:

Let us return for a moment to continuous gray levels r and z (considered continuous random variables), and let $p_r(r)$ and $p_z(z)$ denote their corresponding continuos probability density functions. In this notation, r and z denote the gray levels of the input and output (processed) images, respectively.We can estimate $p_r(r)$ from the given input image, while $p_z(z)$ is the specified probability density function that we wish the output image to have.

Let s be a random variable with the property

$$s = T(r) = \int_0^r p_r(w)\, dw$$

where w is a dummy variable of integration.We recognize this expression as the continuos version of histogram equalization. Suppose next that we define a random variable z with the property

$$G(z) = \int_0^z p_z(t)\, dt = s$$

where t is a dummy variable of integration. It then follows from these two equations that G(z)=T(r) and, therefore, that z must satisfy the condition

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

The transformation T(r) can be obtained once $p_r(r)$ has been estimated from the input image. Similarly, the transformation function G(z) can be obtained because $p_z(z)$ is given. Assuming that G^{-1} exists and that it satisfies conditions (a) and (b) in the histogram equalization process, the above three equations show that an image with a specified probability density function can be obtained from an input image by using the following procedure:

(1) Obtain the transformation function T(r).

(2) To obtain the transformation function G(z).

(3) Obtain the inverse transformation function G^{-1}

(4) Obtain the output image by applying above Eq. to all the pixels in the input image.

The result of this procedure will be an image whose gray levels, z, have the specified probability density function $p_z(z)$. Although the procedure just described is straightforward in principle, it is seldom possible in practice to obtain analytical expressions for T(r) and for G^{-1}. Fortunately, this problem is simplified considerably in the case of discrete values.The price we pay is the same as in histogram equalization,where only an approximation to the desired histogram is achievable. In spite of this, however, some very useful results can be obtained even with crude approximations.

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

$$= \sum_{j=0}^{k} \frac{n_j}{n} \qquad k = 0, 1, 2, \ldots, L-1$$

where n is the total number of pixels in the image, n_j is the number of pixels with gray level r_j, and L is the number of discrete gray levels. Similarly, the discrete formulation is obtained from the given histogram $p_z(z_i)$, i=0, 1, 2,……, L-1, and has the form

$$v_k = G(z_k) = \sum_{i=0}^{k} p_z(z_i) = s_k \qquad k = 0, 1, 2, \ldots, L-1.$$

As in the continuos case, we are seeking values of z that satisfy this equation. The variable v_k was added here for clarity in the discussion that follows. Finally, the discrete version of the above Eqn. is given by

$$z_k = G^{-1}[T(r_k)] \qquad k = 0, 1, 2, \ldots, L-1$$

Or

$$z_k = G^{-1}(s_k) \qquad k = 0, 1, 2, \ldots, L-1.$$

Implementation:

We start by noting the following: (1) Each set of gray levels $\{r_j\}$, $\{s_j\}$, and $\{z_j\}$, j=0, 1, 2, p , L-1, is a one-dimensional array of dimension L X 1. (2) All mappings from r to s and from s to z are simple table lookups between a given pixel value and these arrays. (3) Each of the elements of these arrays, for example, s_k, contains two important pieces of information: The subscript k denotes the location of the element in the array, and s denotes the value at that location. (4) We need to be concerned only with integer pixel values. For example, in the case of an 8-bit image, L=256 and the elements of each of the arrays just mentioned are integers between 0 and 255. This implies that we now work with gray level values in the interval [0, L-1] instead of the normalized interval [0, 1] that we used before to simplify the development of histogram processing techniques.

In order to see how histogram matching actually can be implemented, consider Fig. 5(a), ignoring for a moment the connection shown between this figure and Fig. 5(c). Figure 5(a) shows a hypothetical discrete transformation function s=T(r) obtained from a given image. The first gray level in the image, r_1 , maps to s_1 ; the second gray level, r_2 , maps to s_2 ; the kth level r_k maps to s_k; and so on (the important point here is the ordered correspondence between these values). Each value s_j in the array is precomputed, so the process of mapping simply uses the actual value of a pixel as an index in an array to determine the corresponding value of s. This process is particularly easy because we are dealing with integers. For example, the s mapping for an 8-bit pixel with value 127 would be found in the 128th position in array $\{s_j\}$ (recall that we start at 0) out of the possible 256 positions. If we stopped here and mapped the value of each pixel of an input image by the
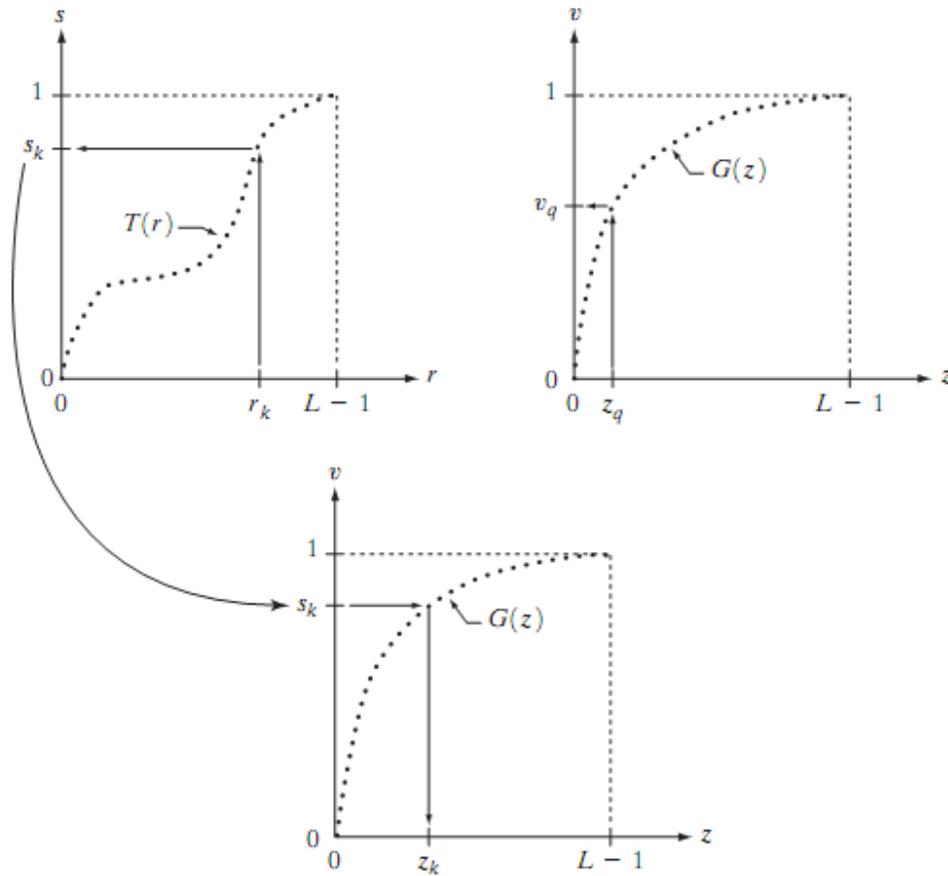
Fig.5. (a) Graphical interpretation of mapping from r_k to s_k via T(r). (b) Mapping of z_q to its corresponding value v_q via G(z) (c) Inverse mapping from s_k to its corresponding value of z_k.

method just described, the output would be a histogram-equalized image. In order to implement histogram matching we have to go one step further. Figure 5(b) is a hypothetical transformation function G obtained from a given histogram $p_z(z)$. For any z_q , this transformation function yields a corresponding value v_q. This mapping is shown by the arrows in Fig. 5(b). Conversely, given any value v_q, we would find the corresponding value z_q from G^{-1}. In terms of the figure, all this means graphically is that we would reverse the direction of the arrows to map v_q into its corresponding z_q. However, we know from the definition that v=s for corresponding subscripts, so we can use exactly this process to find the z_k corresponding to any value s_k that we computed previously from the equation $s_k = T(r_k)$ .This idea is shown in Fig.5(c).

Since we really do not have the z's (recall that finding these values is precisely the objective of histogram matching),we must resort to some sort of iterative scheme to find z from s.The fact that we are dealing with integers makes this a particularly simple process. Basically, because $v_k = s_k$, we have that the z's for which we are looking must satisfy the equation $G(z_k)=s_k$, or $(G(z_k)-s_k)=0$. Thus, all we have to do to find the value of z_k corresponding to s_k is to iterate on values of z such that this equation is satisfied for k=0,1,2,…....., L-1. We do not have to find the inverse of G because we are going to iterate on z. Since we are dealing with integers, the closest we can get to satisfying the equation $(G(z_k)-s_k)=0$ is to let $z_k= \hat{z}$ for each value of k, where $\hat{z}$ is the smallest integer in the interval [0, L-1] such that

$$\left(G(\hat{z}) - s_k\right) \geq 0 \qquad k = 0, 1, 2, \dots, L - 1.$$

Given a value s_k, all this means conceptually in terms of Fig. 5(c) is that we would start with and increase it in integer steps until Eq is satisfied, at which point we let repeating this process for all values of k would yield all the required mappings from s to z, which constitutes the implementation of Eq. In practice, we would not have to start with each time because the values of sk are known to increase monotonically. Thus, for k=k+1, we would start with $\hat{z} = z_k$ and increment in integer values from there.

6. Write about Local enhancement.

Local Enhancement:

The histogram processing methods discussed in the previous two sections are global, in the sense that pixels are modified by a transformation function based on the gray-level content of an entire image. Although this global approach is suitable for overall enhancement, there are cases in which it is necessary to enhance details over small areas in an image. The number of pixels in these areas may have negligible influence on the computation of a global transformation whose shape does not necessarily guarantee the desired local enhancement. The solution is to devise transformation functions based on the gray-level distribution—or other properties—in the neighborhood of every pixel in the image.

The histogram processing techniques are easily adaptable to local enhancement.The procedure is to define a square or rectangular neighborhood and move the center of this area from pixel to pixel. At each location, the histogram of the points in the neighborhood is computed and either a histogram equalization or histogram specification transformation function is obtained. This function is finally used to map the gray level of the

pixel centered in the neighborhood.The center of the neighborhood region is then moved to an adjacent pixel location and the procedure is repeated. Since only one new row or column of the neighborhood changes during a pixel-to-pixel translation of the region, updating the histogram obtained in the previous location with the new data introduced at each motion step is possible. This approach has obvious a dvantages over repeatedly computing the histogram over all pixels in the neighborhood region each time the region is moved one pixel location.Another approach used some times to reduce computation is to utilize nonoverlapping regions, but this method usually produces an undesirable checkerboard effect.

7. What is meant by image subtaction? Discuss various areas of application of image subtraction.

Image Subtraction:

The difference between two images f(x, y) and h(x, y), expressed as

$$g(x, y) = f(x, y) - h(x, y),$$

is obtained by computing the difference between all pairs of corresponding pixels from f and h. The key usefulness of subtraction is the enhancement of differences between images. The higher-order bit planes of an image carry a significant amount of visually relevant detail, while the lower planes contributemore to fine (often imperceptible) detail. Figure 7(a) shows the fractal image used earlier to illustrate the concept of bit planes. Figure 7(b) shows the result of discarding (setting to zero) the four least significant bit planes of the original image.The images are nearly identical visually, with the exception of a very slight drop in overall contrast due to less variability of the graylevel values in the image of Fig. 7(b).The pixel-by-pixel difference between these two images is shown in Fig. 7(c).The differences in pixel values are so small that the difference image appears nearly black when displayed on an 8-bit display. In order to bring out more detail,we can perform a contrast stretching transformation. We chose histogram equalization, but an appropriate power-law transformation would have done the job also. The result is shown in Fig. 7(d). This is a very useful image for evaluating the effect of setting to zero the lower-order planes.
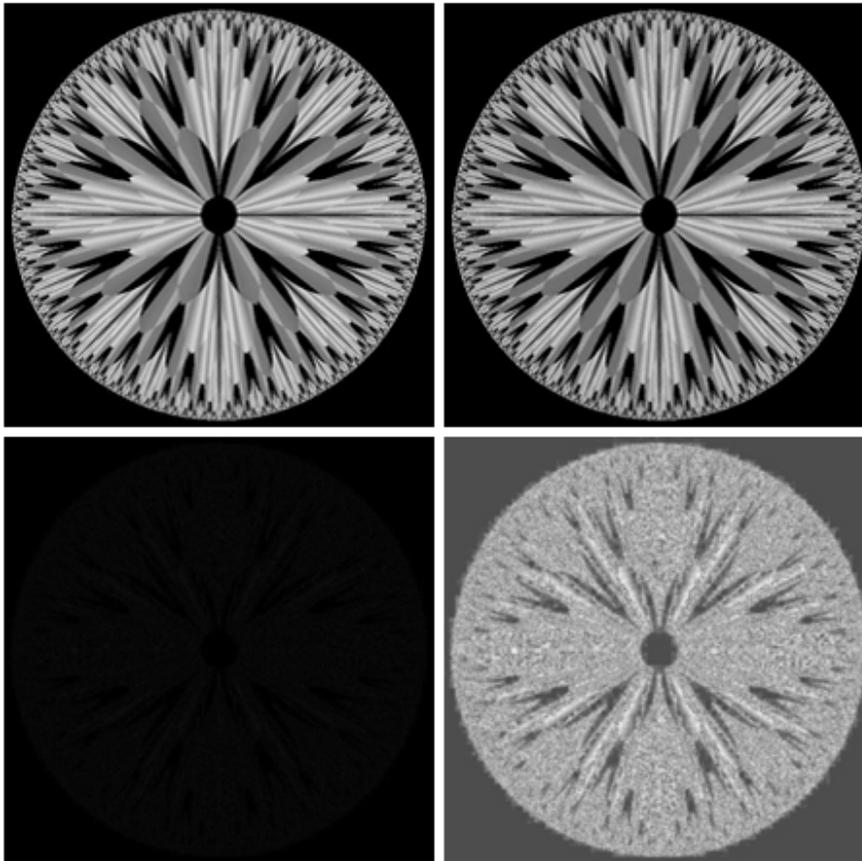
Fig.7 (a) Original fractal image (b) Result of setting the four lower-order bit planes to zero (c) Difference between (a) and(b)   (d) Histogram equalized difference image.

One of the most commercially successful and beneficial uses of image subtraction is in the area of medical imaging called mask mode radiography. In this case h(x, y), the mask, is an X-ray image of a region of a patient's body captured by an intensified TV camera (instead of traditional X-ray film) located opposite an X-ray source.The procedure consists of injecting a contrast medium into the patient's bloodstream, taking a series of images of the same anatomical region as h(x, y), and subtracting this mask from the series of incoming images after injection of the contrast medium. The net effect of subtracting the mask from each sample in the incoming stream of TV images is that the areas that are different between f(x, y) and h(x, y) appear in the output image as enhanced detail. Because images can be captured at TV rates, this procedure in essence gives a movie showing how the contrast medium propagates through the various arteries in the area being observed.

8. Explain about image averaging process.

Image Averaging:

Consider a noisy image g(x, y) formed by the addition of noise h(x, y) to an original image f(x,y); that is,

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

where the assumption is that at every pair of coordinates (x, y) the noise is uncorrelated and has zero average value. The objective of the following procedure is to reduce the noise content by adding a set of noisy images, $\{g_i(x, y)\}$. If the noise satisfies the constraints just stated, it can be shown that if an image $\bar{g}(x, y)$ is formed by averaging K different noisy images,

$$\bar{g}(x, y) = \frac{1}{K} \sum_{i=1}^{K} g_i(x, y)$$

Then it follows that

$$E\{\bar{g}(x, y)\} = f(x, y)$$

and

$$\sigma^2_{\bar{g}(x,y)} = \frac{1}{K} \sigma^2_{\eta(x,y)}$$

Where $E\{\bar{g}(x, y)\}$ is the expected value of $\bar{g}$ and $\sigma^2_{\bar{g}(x,y)}$ and $\sigma^2_{\bar{\eta}(x,y)}$ are the variances of $\bar{g}$ and η, all at coordinates (x, y). The standard deviation at any point in the average image is

$$\sigma_{\bar{g}(x,y)} = \frac{1}{\sqrt{K}} \sigma_{\eta(x,y)}.$$

As K increases, the above equations indicate that the variability (noise) of the pixel values at each location (x, y) decreases.Because $E\{\bar{g}(x, y)\} = f(x, y),$ this means that $\bar{g}(x, y)$ approaches f(x, y) as the number of noisy images used in the averaging process increases. In practice, the images $g_i(x, y)$ must be registered (aligned) in order to avoid the introduction of blurring and other artifacts in the output image.

9. Discuss about the mechanics of filtering in spatial domain. Mention the points to be considered in implementation neighbourhood operations for spatial filtering.

Basics of Spatial Filtering:

Some neighborhood operations work with the values of the image pixels in the neighborhood and the corresponding values of a subimage that has the same dimensions as the neighborhood.The subimage is called a filter,mask, kernel, template, or window,with the first three terms being the most prevalent terminology.The values in a filter subimage are referred to as coefficients, rather than pixels. The concept of filtering has its roots in the use of the Fourier transform for signal processing in the so-called frequency domain. We use the term spatial filtering to differentiate this type of process from the more traditional frequency domain filtering.

The mechanics of spatial filtering are illustrated in Fig.9.1. The process consists simply of moving the filter mask from point to point in an image. At each point (x, y), the response of the filter at that point is calculated using a predefined relationship. The response is given by a sum of products of the filter coefficients and the corresponding image pixels in the area spanned by the filter mask. For the 3 x 3 mask shown in Fig. 9.1, the result (or response), R, of linear filtering with the filter mask at a point (x, y) in the image is

$$R = w(-1,-1)f(x - 1, y - 1) + w(-1,0)f(x - 1, y) + \cdots$$
$$+ w(0,0)f(x, y) + \cdots + w(1,0)f(x + 1, y) + w(1,1)f(x + 1, y + 1),$$

which we see is the sum of products of the mask coefficients with the corresponding pixels directly under the mask. Note in particular that the coefficient w(0, 0) coincides with image

value f(x, y), indicating that the mask is centered at (x, y) when the computation of the sum of products takes place. For a mask of size m x n,we assume that m=2a+1 and n=2b+1,where a and b are nonnegative integers.
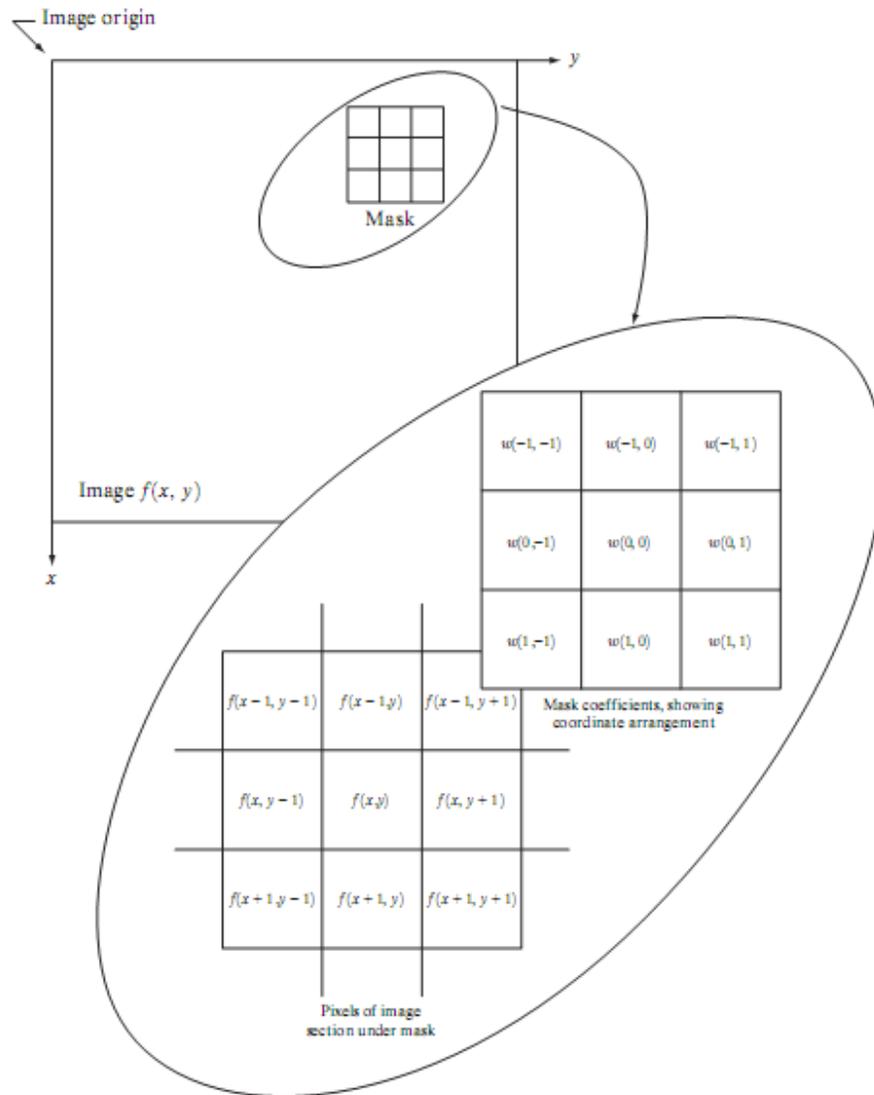
Fig.9.1 The mechanics of spatial filtering. The magnified drawing shows a 3X3 mask and the image section directly under it; the image section is shown displaced out from under the mask for ease of readability.

In general, linear filtering of an image f of size M x N with a filter mask of size m x n is given by the expression:

$$g(x, y) = \sum_{s=-a}^{a} \sum_{t=-b}^{b} w(s, t)f(x + s, y + t)$$

where, from the previous paragraph, a=(m-1)/2 and b=(n-1)/2. To generate a complete filtered image this equation must be applied for x=0,1,2,......, M-1 and y=0,1,2,......, N-1. In this way, we are assured that the mask processes all pixels in the image. It is easily verified when m=n=3 that this expression reduces to the example given in the previous paragraph.

 The process of linear filtering is similar to a frequency domain concept called convolution. For this reason, linear spatial filtering often is referred to as "convolving a mask with an image." Similarly, filter masks are sometimes called convolution masks. The term convolution kernel also is in common use. When interest lies on the response, R, of an m x n mask at any point (x,y), and not on the mechanics of implementing mask convolution, it is common practice to simplify the notation by using the following expression:

$$R = w_1 z_1 + w_2 z_2 + \ldots + w_{mn} z_{mn}$$
$$= \sum_{i=1}^{mn} w_i z_i$$

where the w's are mask coefficients, the z's are the values of the image graylevels corresponding to those coefficients, and mn is the total number of coefficients in the mask. For the 3 x 3 general mask shown in Fig.9.2 the response at any point (x, y) in the image is given by

$$R = w_1 z_1 + w_2 z_2 + \ldots w_9 z_9$$
$$= \sum_{i=1}^{9} w_i z_i.$$

w_1	w_2	w_3
w_4	w_5	w_6
w_7	w_8	w_9

Fig.9.2 Another representation of a general 3 x 3 spatial filter mask.

An important consideration in implementing neighborhood operations for spatial filtering is the issue of what happens when the center of the filter approaches the border of the image.Consider for simplicity a square mask of size n x n.At least one edge of such a mask will coincide with the border of the image when the center of the mask is at a distance of (n-1)/2 pixels away from the border of the image. If the center of the mask moves any closer to the border, one or more rows or columns of the mask will be located outside the image plane.There are several ways to handle this situation.The simplest is to limit the excursions of the center of the mask to be at a distance no less than (n-1)/2 pixels from the border. The resulting filtered image will be smaller than the original, but all the pixels in the filtered imaged will have been processed with the full mask. If the result is required to be the same size as the original, then the approach typically employed is to filter all pixels only with the section of the mask that is fully contained in the image.With this approach, there will be bands of pixels near the border that will have been processed with a partial filter mask.Other approaches include "padding" the image by adding rows and columns of 0's (or other constant gray level), or padding by replicating rows or columns.The padding is then stripped off at the end of the process.

This keeps the size of the filtered image the same as the original, but the values of the padding will have an effect near the edges that becomes more prevalent as the size of the mask increases.The only way to obtain a perfectly filtered result is to accept a somewhat smaller filtered image by limiting the excursions of the center of the filter mask to a distance no less than (n-1)/2 pixels from the border of the original image.

10. Write about Smoothing Spatial filters.

Smoothing Spatial Filters:

Smoothing filters are used for blurring and for noise reduction. Blurring is used in preprocessing steps, such as removal of small details from an image prior to (large) object extraction, and bridging of small gaps in lines or curves. Noise reduction can be accomplished by blurring with a linear filter and also by non-linear filtering.

(1) Smoothing Linear Filters:

The output (response) of a smoothing, linear spatial filter is simply the average of the pixels contained in the neighborhood of the filter mask. These filters sometimes are called averaging filters. The idea behind smoothing filters is straightforward.By replacing the value of every pixel

in an image by the average of the gray levels in the neighborhood defined by the filter mask, this process results in an image with reduced "sharp" transitions in gray levels. Because random noise typically consists of sharp transitions in gray levels, the most obvious application of smoothing is noise reduction.However, edges (which almost always are desirable features of an image) also are characterized by sharp transitions in gray levels, so averaging filters have the undesirable side effect that they blur edges. Another application of this type of process includes the smoothing of false contours that result from using an insufficient number of gray levels.
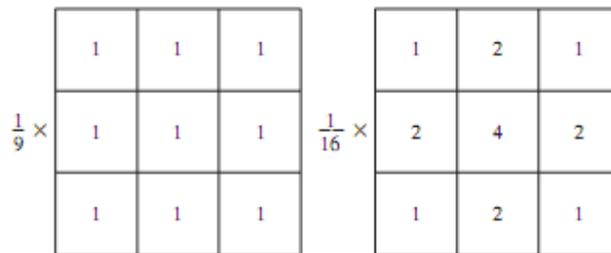
Fig.10.1 Two 3 x 3 smoothing (averaging) filter masks.The constant multiplier in front of each mask is equal to the sum of the values of its coefficients, as is required to compute an average.

A major use of averaging filters is in the reduction of "irrelevant" detail in an image. By "irrelevant"we mean pixel regions that are small with respect to the size of the filter mask.

Figure 10.1 shows two 3 x 3 smoothing filters. Use of the first filter yields the standard average of the pixels under the mask.This can best be seen by substituting the coefficients of the mask in

$$R = \frac{1}{9} \sum_{i=1}^{9} z_i,$$

which is the average of the gray levels of the pixels in the 3 x 3 neighborhood defined by the mask.Note that, instead of being 1/9, the coefficients of the filter are all 1's.The idea here is that it is computationally more efficient to have coefficients valued 1. At the end of the filtering process the entire image is divided by 9. An m x n mask would have a normalizing constant equal to 1/mn.

A spatial averaging filter in which all coefficients are equal is sometimes called a box filter.

The second mask shown in Fig.10.1 is a little more interesting. This mask yields a so-called weighted average, terminology used to indicate that pixels are multiplied by different coefficients, thus giving more importance (weight) to some pixels at the expense of others. In the mask shown in Fig. 10.1(b) the pixel at the center of the mask is multiplied by a higher value than any other, thus giving this pixel more importance in the calculation of the average.The other pixels are inversely weighted as a function of their distance from the center of the mask. The diagonal terms are further away from the center than the orthogonal neighbors (by a factor of √2) and, thus, are weighed less than these immediate neighbors of the center pixel. The basic strategy behind weighing the center point the highest and then reducing the value of the coefficients as a function of increasing distance from the origin is simply an attempt to reduce blurring in the smoothing process. We could have picked other weights to accomplish the same general objective. However, the sum of all the coefficients in the mask of Fig. 10.1(b) is equal to 16, an attractive feature for computer implementation because it has an integer power of 2. In practice, it is difficult in general to see differences between images smoothed by using either of the masks in Fig. 10.1, or similar arrangements, because the area these masks span at any one location in an image is so small.

The general implementation for filtering an M x N image with a weighted averaging filter of size m x n (m and n odd) is given by the expression

$$g(x, y) = \frac{\sum_{s=-a}^{a} \sum_{t=-b}^{b} w(s, t) f(x + s, y + t)}{\sum_{s=-a}^{a} \sum_{t=-b}^{b} w(s, t)}$$

.

(2) Order-Statistics Filters:

Order-statistics filters are nonlinear spatial filters whose response is based on ordering (ranking) the pixels contained in the image area encompassed by the filter, and then replacing the value of the center pixel with the value determined by the ranking result. The best-known example in this category 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 (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 impulse noise, also called salt-and-pepper noise because of its appearance as white and black dots superimposed on an image.

The median, ε, of a set of values is such that half the values in the set are less than or equal to ε, and half are greater than or equal to ε. In order to perform median filtering at a point in an image, we first sort the values of the pixel in question and its neighbors, determine their median, and

assign this value to that pixel. For example, in a 3 x 3 neighborhood the median is the 5th largest value, in a 5 x 5 neighborhood the 13th largest value, and so on. When several values in a neighborhood are the same, all equal values are grouped. For example, suppose that a 3 x 3 neighborhood has values (10, 20, 20, 20, 15, 20, 20, 25, 100). These values are sorted as (10, 15, 20, 20, 20, 20, 20, 25, 100), which results in a median of 20. Thus, the principal function of median filters is to force points with distinct gray levels to be more like their neighbors. In fact, isolated clusters of pixels that are light or dark with respect to their neighbors, and whose area is less than $n^2/2$ (one-half the filter area), are eliminated by an n x n median filter. In this case "eliminated" means forced to the median intensity of the neighbors. Larger clusters are affected considerably less.

11. What is meant by the Gradiant and the Laplacian? Discuss their role in image enhancement.

Use of Second Derivatives for Enhancement–The Laplacian:

The approach basically consists of defining a discrete formulation of the second-order derivative and then constructing a filter mask based on that formulation. We are interested in isotropic filters, whose response is independent of the direction of the discontinuities in the image to which the filter is applied. In other words, isotropic filters are rotation invariant, in the sense that rotating the image and then applying the filter gives the same result as applying the filter to the image first and then rotating the result.

Development of the method:

It can be shown (Rosenfeld and Kak [1982]) that the simplest isotropic derivative operator is the Laplacian, which, for a function (image) f(x, y) of two variables, is defined as

$$\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}.$$

Because derivatives of any order are linear operations, the Laplacian is a linear operator. In order to be useful for digital image processing, this equation needs to be expressed in discrete form. There are several ways to define a digital Laplacian using neighborhoods. digital second.Taking into account that we now have two variables, we use the following notation for the partial second-order derivative in the x-direction:

$$\frac{\partial^2 f}{\partial^2 x^2} = f(x + 1, y) + f(x - 1, y) - 2f(x, y)$$

and, similarly in the y-direction, as

$$\frac{\partial^2 f}{\partial^2 y^2} = f(x, y + 1) + f(x, y - 1) - 2f(x, y)$$

The digital implementation of the two-dimensional Laplacian in Eq. is obtained by summing these two components

$$\nabla^2 f = \left[f(x + 1, y) + f(x - 1, y) + f(x, y + 1) + f(x, y - 1) \right] - 4f(x, y). \qquad ($$

This equation can be implemented using the mask shown in Fig.11.1(a), which gives an isotropic result for rotations in increments of 90°.

The diagonal directions can be incorporated in the definition of the digital Laplacian by adding two more terms to Eq., one for each of the two diagonal directions.The form of each new term is the same as either Eq.

0	1	0
1	−4	1
0	1	0

1	1	1
1	−8	1
1	1	1

0	−1	0
−1	4	−1
0	−1	0

−1	−1	−1
−1	8	−1
−1	−1	−1

a	b
c	d

Fig.11.1. (a) Filter mask used to implement the digital Laplacian (b) Mask used to implement an extension of this equation that includes the diagonal neighbors. (c) and (d) Two other implementations of the Laplacian.

but the coordinates are along the diagonals. Since each diagonal term also contains a −2f(x, y) term, the total subtracted from the difference terms now would be −8f(x, y). The mask used to implement this new definition is shown in Fig.11.1(b). This mask yields isotropic results for increments of 45°. The other two masks shown in Fig. 11 also are used frequently in practice.

They are based on a definition of the Laplacian that is the negative of the one we used here. As such, they yield equivalent results, but the difference in sign must be kept in mind when combining (by addition or subtraction) a Laplacian-filtered image with another image.

Because the Laplacian is a derivative operator, its use highlights gray-level discontinuities in an image and deemphasizes regions with slowly varying gray levels.This will tend to produce images that have grayish edge lines and other discontinuities, all superimposed on a dark, featureless background.Background features can be "recovered" while still preserving the sharpening effect of the Laplacian operation simply by adding the original and Laplacian images. As noted in the previous paragraph, it is important to keep in mind which definition of the Laplacian is used. If the definition used has a negative center coefficient, then we subtract, rather

than add, the Laplacian image to obtain a sharpened result.Thus, the basic way in which we use the Laplacian for image enhancement is as follows:

$$g(x, y) = \begin{cases} f(x, y) - \nabla^2 f(x, y) & \text{if the center coefficient of the Laplacian mask is negative} \\ f(x, y) + \nabla^2 f(x, y) & \text{if the center coefficient of the Laplacian mask is positive.} \end{cases}$$

Use of First Derivatives for Enhancement—The Gradient:

First derivatives in image processing are implemented using the magnitude of the gradient. For a function f(x, y), the gradient of f at coordinates (x, y) is defined as the two-dimensional column vector

$$\nabla \mathbf{f} = \begin{bmatrix} G_x \\ G_y \end{bmatrix} = \begin{bmatrix} \dfrac{\partial f}{\partial x} \\ \dfrac{\partial f}{\partial y} \end{bmatrix}.$$

The magnitude of this vector is given by

$$\begin{aligned} \nabla f &= \text{mag}(\nabla \mathbf{f}) \\ &= \left[G_x^2 + G_y^2 \right]^{1/2} \\ &= \left[\left(\frac{\partial f}{\partial x} \right)^2 + \left(\frac{\partial f}{\partial y} \right)^2 \right]^{1/2}. \end{aligned}$$

The components of the gradient vector itself are linear operators, but the magnitude of this vector obviously is not because of the squaring and square root operations. On the other hand, the partial derivatives are not rotation invariant (isotropic), but the magnitude of the gradient vector is. Although it is not strictly correct, the magnitude of the gradient vector often is referred to as the gradient.

The computational burden of implementing over an entire image is not trivial, and it is common practice to approximate the magnitude of the gradient by using absolute values instead of squares and square roots:

$$\nabla f \approx |G_x| + |G_y|.$$

This equation is simpler to compute and it still preserves relative changes in gray levels, but the isotropic feature property is lost in general. However, as in the case of the Laplacian, the isotropic properties of the digital gradient defined in the following paragraph are preserved only for a limited number of rotational increments that depend on the masks used to approximate the derivatives. As it turns out, the most popular masks used to approximate the gradient give the same result only for vertical and horizontal edges and thus the isotropic properties of the gradient are preserved only for multiples of 90°.

As in the case of the Laplacian, we now define digital approximations to the preceding equations, and from there formulate the appropriate filter masks. In order to simplify the discussion that follows, we will use the notation in Fig. 11.2 (a) to denote image points in a 3 x 3 region. For example, the center point, z5 , denotes f(x, y), z1 denotes f(x-1, y-1), and so on. The simplest approximations to a first-order derivative that satisfy the conditions stated in that section are $G_x = (z_8 - z_5)$ and $G_y = (z_6 - z_5)$. Two other definitions proposed by Roberts [1965] in the early development of digital image processing use cross differences:

$$G_x = (z_9 - z_5) \quad \text{and} \quad G_y = (z_8 - z_6).$$

we compute the gradient as

$$\nabla f = [(z_9 - z_5)^2 + (z_8 - z_6)^2]^{1/2}$$

If we use absolute values, then substituting the quantities in the equations gives us the following approximation to the gradient:

$$\nabla f \approx |z_9 - z_5| + |z_8 - z_6|.$$

This equation can be implemented with the two masks shown in Figs. 11.2 (b) and(c). These masks are referred to as the Roberts cross-gradient operators. Masks of even size are awkward to implement. The smallest filter mask in which we are interested is of size 3 x 3.An approximation using absolute values, still at point z5 , but using a 3*3 mask, is

$$\nabla f \approx |(z_7 + 2z_8 + z_9) - (z_1 + 2z_2 + z_3)| + |(z_3 + 2z_6 + z_9) - (z_1 + 2z_4 + z_7)|.$$

The difference between the third and first rows of the 3 x 3 image region approximates the derivative in the x-direction, and the difference between the third and first columns approximates the derivative in the y-direction. The masks shown in Figs. 11.2 (d) and (e), called the Sobel operators. The idea behind using a weight value of 2 is to achieve some smoothing by giving more importance to the center point. Note that the coefficients in all the masks shown in Fig. 11.2 sum to 0, indicating that they would give a response of 0 in an area of constant gray level, as expected of a derivative operator.
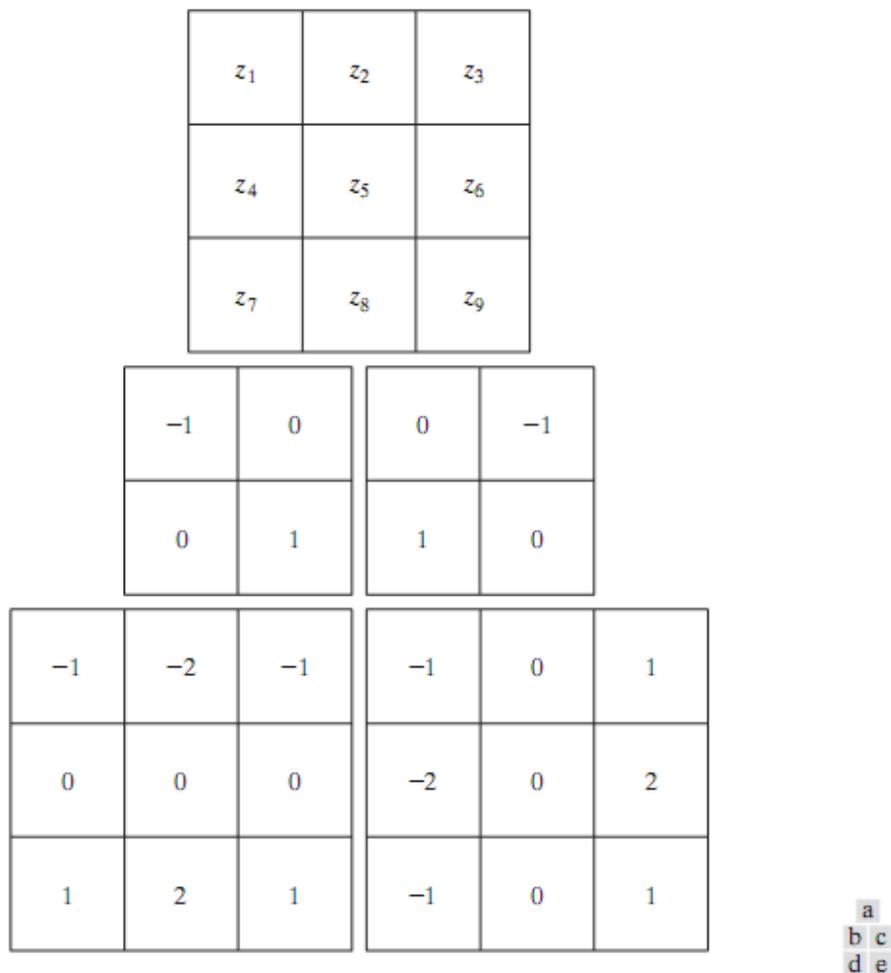
Fig.11.2 A 3 x 3 region of an image (the z's are gray-level values) and masks used to compute the gradient at point labeled z_5 . All masks coefficients sum to zero, as expected of a derivative operator.

1.Discuss the frequency domain techniques of image enhancement in detail.

 Enhancement In Frequency Domain:

The frequency domain methods of image enhancement are based on convolution theorem. This is represented as,

$$g(x, y) = h(x, y)*f(x, y)$$

Where.

$$g(x, y) = \text{Resultant image}$$
$$h(x, y) = \text{Position invariant operator}$$
$$f(x, y) = \text{Input image}$$

The Fourier transform representation of equation above is,

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

The function H (u, v) in equation is called transfer function. It is used to boost the edges of input image f (x, y) to emphasize the high frequency components.

The different frequency domain methods for image enhancement are as follows.

1. Contrast stretching.
2. Clipping and thresholding.
3. Digital negative.
4. Intensity level slicing and
5. Bit extraction.

1. Contrast Stretching:

 Due to non-uniform lighting conditions, there may be poor contrast between the background and the feature of interest. Figure 1.1 (a) shows the contrast stretching transformations.
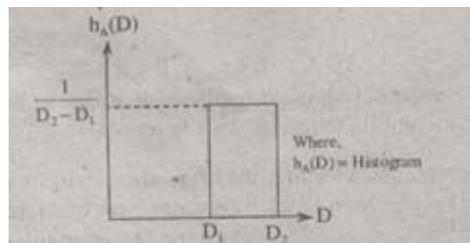
Fig.1.1 (a) Histogram of input image
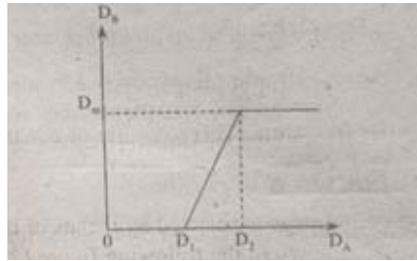
Fig.1.1 (b) Linear Law
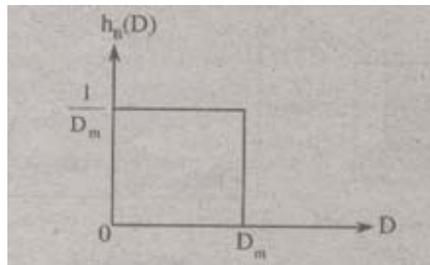
Fig.1.1 (c) Histogram of the transformed image

These stretching transformations are expressed as

In the area of stretching the slope of transformation is considered to be greater than unity. The parameters of stretching transformations i.e., a and b can be determined by examining the histogram of the image.

2. Clipping and Thresholding:

Clipping is considered as the special scenario of contrast stretching. It is the case in which the parameters are $\alpha = \gamma = 0$. Clipping is more advantageous for reduction of noise in input signals of range [a, b].

Threshold of an image is selected by means of its histogram. Let us take the image shown in the following figure 1.2.
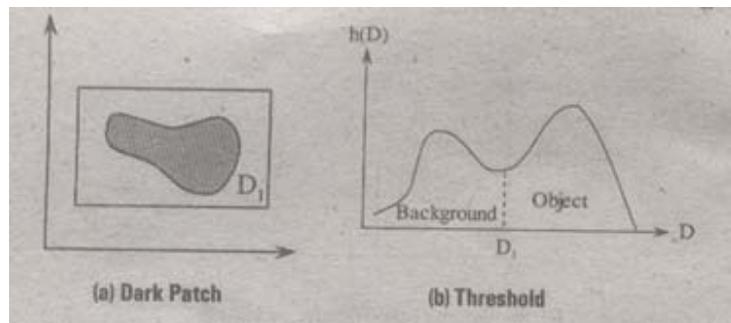
Fig. 1.2

The figure 1.2 (b) consists of two peaks i.e., background and object. At the abscissa of histogram minimum (D1) the threshold is selected. This selected threshold (D1) can separate background and object to convert the image into its respective binary form. The thresholding transformations are shown in figure 1.3.
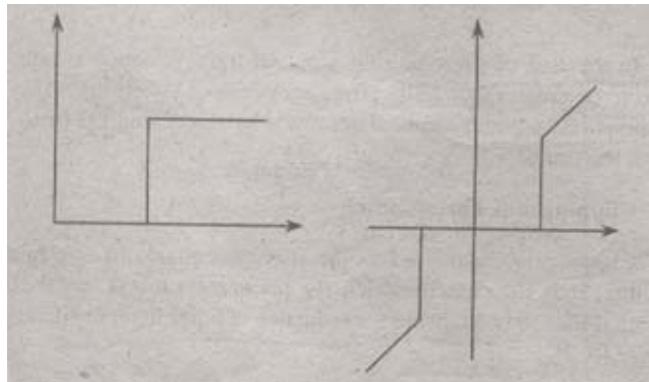
Fig.1.3

3. Digital Negative:

The digital negative of an image is achieved by reverse scaling of its grey levels to the transformation. They are much essential in displaying of medical images.

A digital negative transformation of an image is shown in figure 1.4.
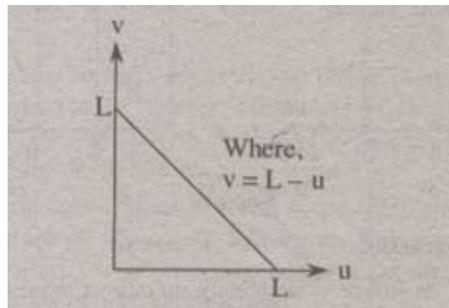
Fig.1.4

4. Intensity Level Slicing:

The images which consist of grey levels in between intensity at background and other objects require to reduce the intensity of the object. This process of changing intensity level is done with the help of intensity level slicing. They are expressed as

$$V = \begin{cases} L, & a \leq u \leq b \\ 0, & \text{elsewhere} \end{cases} \quad \text{without background}$$

$$\text{And} \quad V = \begin{cases} L, & a \leq u \leq b \\ u^c, & \text{elsewhere} \end{cases} \quad \text{with background}$$

The histogram of input image and its respective intensity level slicing is shown in the figure 1.5.
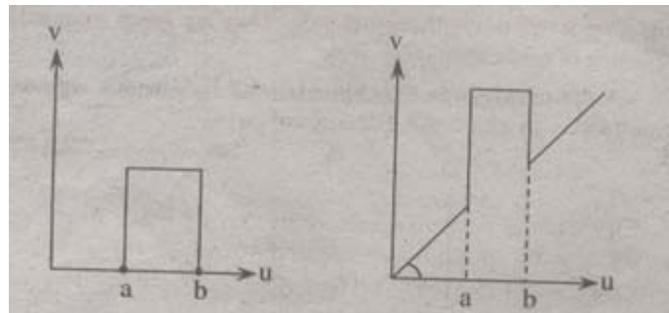
Fig.1.5

When an image is uniformly quantized then, the n[th] most significant bit can be extracted and displayed.

$$\text{Let, } u = k_1 \, 2^{B-1} + k_2 \, 2^{B-2} + \ldots\ldots\ldots\ldots + k_{B-1} \, 2 + k_B$$

Then, the output is expressed as

$$V = \begin{cases} L, & \text{for } k_n = 1 \\ 0, & \text{elsewhere} \end{cases}$$

2. Distinguish between spatial domain and frequency domain enhancement techniques.

The spatial domain refers to the image plane itself, and approaches in this category are based on direct manipulation of pixels in an image. Frequency domain processing techniques are based on modifying the Fourier transform of an image.

The term spatial domain refers to the aggregate of pixels composing an image and spatial domain methods are procedures that operate directly on these pixels. Image processing function in the spatial domain may he expressed as.

$$g(x, y) = T[f(x, y)]$$

Where

f(x, y) is the input image
g(x, y) is the processed image and
T is the operator on f defined over some neighborhood values of (x, y).

Frequency domain techniques are based on convolution theorem. Let g(x, y) be the image formed by the convolution of an image f(x, y) and linear position invariant operation h(x, y) i.e.,

$$g(x, y) = h(x, y) * f(x, y)$$

Applying convolution theorem

$$G(u, v) = H(u, v) \, F(u, v)$$

Where G, H and F are the Fourier transforms of g, h and f respectively. In the terminology of linear system the transform H (u, v) is called the transfer function of the process. The edges in f(x, y) can he boosted by using H (u, v) to emphasize the high frequency components of F (u, v).

3. Explain about Ideal Low Pass Filter (ILPF) in frequency domain.

Lowpass Filter:

The edges and other sharp transitions (such as noise) in the gray levels of an image contribute significantly to the high-frequency content of its Fourier transform. Hence blurring (smoothing) is achieved in the frequency domain by attenuating us the transform of a given image.

$$G\ (u,\ v) = H\ (u,\ v)\ F(u,\ v)$$

where F (u, v) is the Fourier transform of an image to be smoothed. The problem is to select a filter transfer function H (u, v) that yields G (u, v) by attenuating the high-frequency components of F (u, v). The inverse transform then will yield the desired smoothed image g (x, y).

Ideal Filter:

A 2-D ideal lowpass filter (ILPF) is one whose transfer function satisfies the relation

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

where D is a specified nonnegative quantity, and D(u, v) is the distance from point (u, v) to the origin of the frequency plane; that is,

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

Figure 3 (a) shows a 3-D perspective plot of H (u, v) u a function of u and v. The name ideal filter indicates that oil frequencies inside a circle of radius
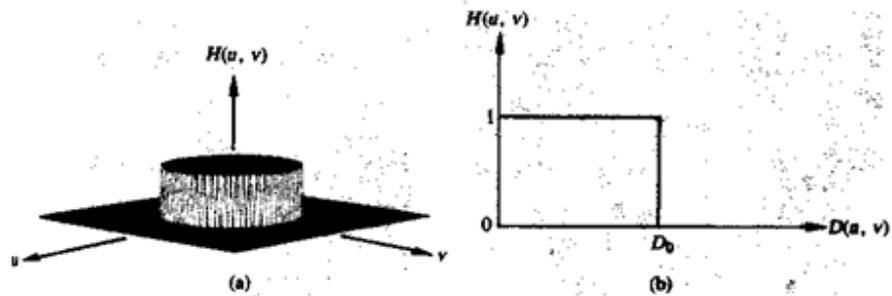
Fig.3 (a) Perspective plot of an ideal lowpass filter transfer function; (b) filter cross section.

Do are passed with no attenuation, whereas all frequencies outside this circle are completely attenuated.

The lowpass filters are radially symmetric about the origin. For this type of filter, specifying a cross section extending as a function of distance from the origin along a radial line is sufficient, as Fig. 3 (b) shows. The complete filter transfer function can then be generated by rotating the cross section 360 about the origin. Specification of radially symmetric filters centered on the N x N frequency square is based on the assumption that the origin of the Fourier transform has been centered on the square.

For an ideal lowpass filter cross section, the point of transition between H(u, v) = 1 and H(u, v) = 0 is often called the cutoff frequency. In the case of Fig.3 (b), for example, the cutoff frequency is Do. As the cross section is rotated about the origin, the point Do traces a circle giving a locus of cutoff frequencies, all of which are a distance Do from the origin. The cutoff frequency concept is quite useful in specifying filter characteristics. It also serves as a common base for comparing the behavior of different types of filters.

The sharp cutoff frequencies of an ideal lowpass filter cannot be realized with electronic components, although they can certainly be simulated in a computer.

4. Discuss about Butterworth lowpass filter with a suitable example.

Butterworth filter:

The transfer function of the Butterworth lowpass (BLPF) of order n and with cutoff frequency locus at a distance Do, from the origin is defined by the relation

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

A perspective plot and cross section of the BLPF function are shown in figure 4.
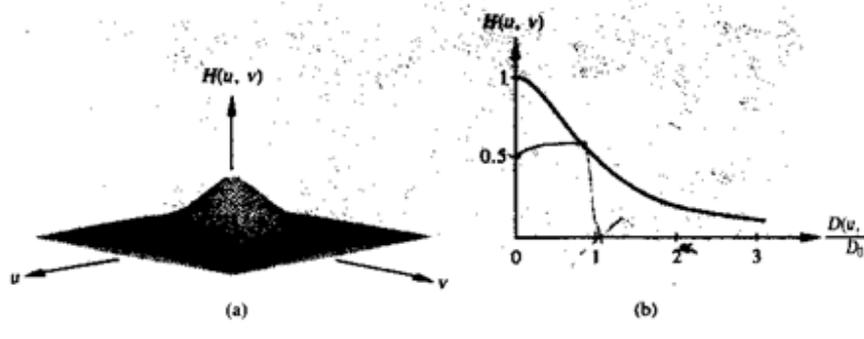
Fig.4 (a) A Butterworth lowpass filter (b) radial cross section for n = 1.

Unlike the ILPF, the BLPF transfer function does not have a sharp discontinuity that establishes a clear cutoff between passed and filtered frequencies. For filters with smooth transfer functions, defining a cutoff frequency locus at points for which H (u, v) is down to a certain fraction of its maximum value is customary. In the case of above Eq. H (u, v) = 0.5 (down 50 percent from its maximum value of 1) when D (u, v) = Do. Another value commonly used is $1/\sqrt{2}$ of the maximum value of H (u, v). The following simple modification yields the desired value when D (u, v) = Do:

$$H(u, v) = \frac{1}{1 + [\sqrt{2} - 1][D(u, v)/D_0]^{2n}}$$
$$= \frac{1}{1 + 0.414[D(u, v)/D_0]^{2n}}$$

5. Discuss about Ideal High Pass Filter and Butterworth High Pass filter.

High pass Filtering:

 An image can be blurred by attenuating the high-frequency components of its Fourier transform. Because edges and other abrupt changes in gray levels are associated with high-frequency components, image sharpening can be achieved in the frequency domain by a high pass filtering process, which attenuates the low-frequency components without disturbing high-frequency information in the Fourier transform.

Ideal filter:

2-D ideal high pass filter (IHPF) is one whose transfer function satisfies the relation

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

where Do is the cutoff distance measured from the origin of the frequency plane. Figure 5.1 shows a perspective plot and cross section of the IHPF function. This filter is the opposite of the ideal lowpass filter, because it completely attenuates all frequencies inside a circle of radius Do while passing, without attenuation, all frequencies outside the circle. As in the case of the ideal lowpass filler, the IHPF is not physically realizable.
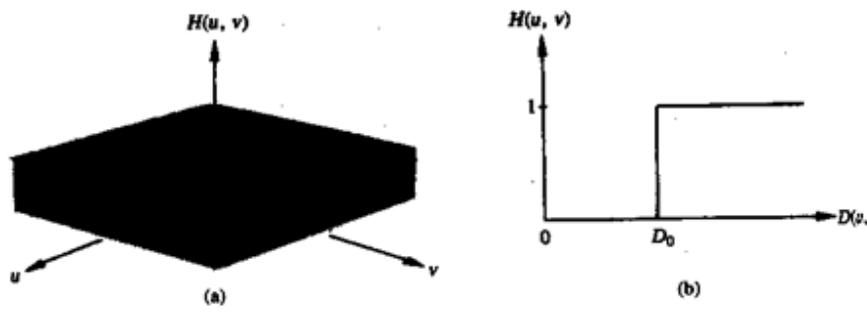
Fig.5.1 Perspective plot and radial cross section of ideal high pass filter

Butterworth filter:

The transfer function of the Butterworth high pass filter (BHPF) of order n and with cutoff frequency locus at a distance Do from the origin is defined by the relation

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

Figure 5.2 shows a perspective plot and cross section of the BHPF function. Note that when D (u, v) = Do, H (u, v) is down to ½ of its maximum value. As in the case of the Butterworth lowpass filter, common practice is to select the cutoff frequency locus at points for which H (u, v) is down to $1/\sqrt{2}$ of its maximum value.

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

$$= \frac{1}{1 + 0.414[D_0/D(u, v)]^{2n}}.$$

(a)   (b)
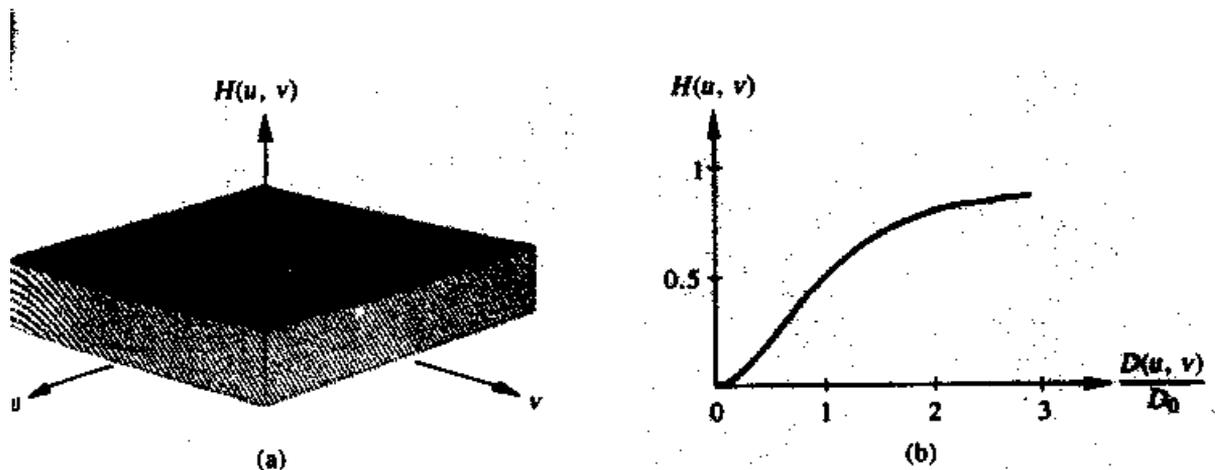
Fig.5.2 Perspective plot and radial cross section for Butterworth High Pass Filter with n = 1

6. Discuss about Gaussian High Pass and Gaussian Low Pass Filter.

Gaussian Lowpass Filters:

The form of these filters in two dimensions is given by

$$H(u, v) = e^{-D^2(u, v)/2\sigma^2}$$

where, D(u, v) is the distance from the origin of the Fourier transform.

a b c
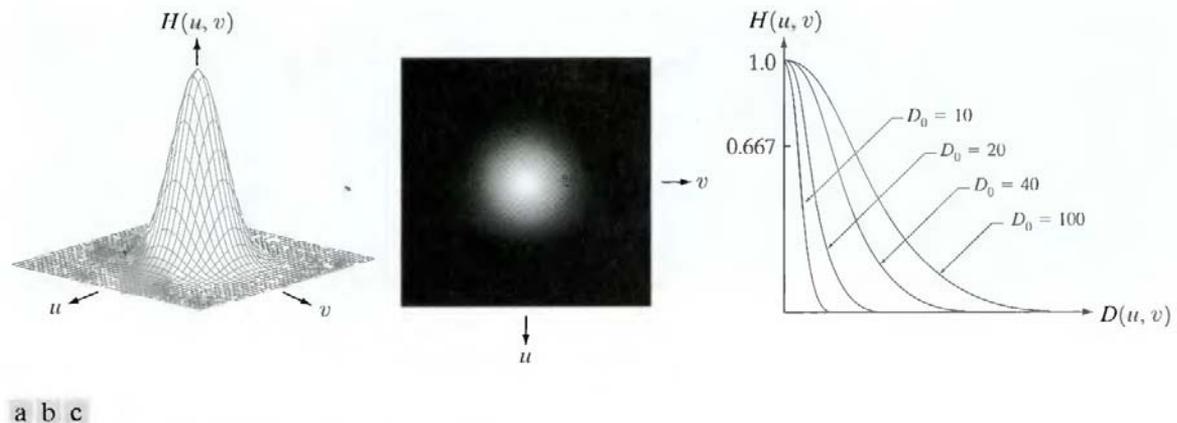
Fig.6.1 (a) Perspective plot of a GLPF transfer function, (b) Filter displayed as an image, (c) Filter radial cross sections for various values of Do.

σ is a measure of the spread of the Gaussian curve. By letting σ = Du, we can express the filter in a more familiar form in terms of the notation:

$$H(u, v) = e^{-D^2(u, v)/2D_0^2}$$

where Do is the cutoff frequency. When D (u, v) = Do, the filter is down to 0.607 of its maximum value.

Gaussian Highpass Filters:

The transfer function of the Gaussian highpass filter (GHPF) with cutoff frequency locus at a distance Do from the origin is given by

$$H(u, v) = 1 - e^{-D^2(u, v)/2D_0^2}$$

The figure 6.2 shows a perspective plot, image, and cross section of the GHPF function.
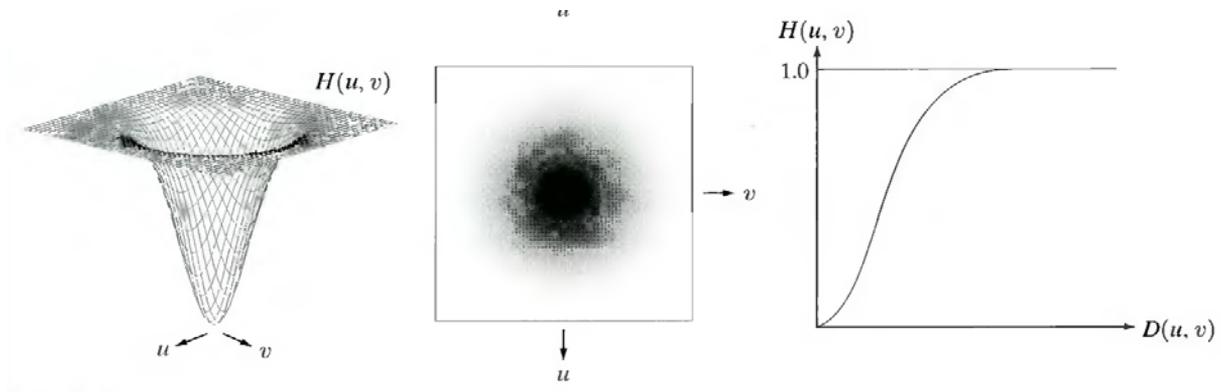
Fig.6.2. Perspective plot, image representation, and cross section of a typical Gaussian high pass filter

Even the filtering of the smaller objects and thin bars is cleaner with the Gaussian filler.

7. Explain how Laplacian is implemented in frequency domain.

The Laplacian in the Frequency Domain:

It can be shown that

$$\Im\left[\frac{d^n f(x)}{dx^n}\right] = (ju)^n F(u).$$

From this simple expression, it follows that

$$\Im\left[\frac{\partial^2 f(x, y)}{\partial x^2} + \frac{\partial^2 f(x, y)}{\partial y^2}\right] = (ju)^2 F(u, v) + (jv)^2 F(u, v)$$
$$= -(u^2 + v^2)F(u, v).$$

The expression inside the brackets on the left side of the above Eq. is recognized as the Laplacian of f(x, y). Thus, we have the important result

$$\Im\left[\nabla^2 f(x, y)\right] = -(u^2 + v^2)F(u, v),$$

which simply says that the Laplacian can be implemented in the frequency domain by using the filter

$$H(u, v) = -(u^2 + v^2).$$

As in all filtering operations, the assumption is that the origin of F (u, v) has been centered by performing the operation f(x, y) (-1) $^{x+y}$ prior to taking the transform of the image. If f (and F) are of size M X N, this operation shifts the center transform so that (u, v) = (0, 0) is at point (M/2, N/2) in the frequency rectangle. As before, the center of the filter function also needs to be shifted:

$$H(u, v) = -\left[(u - M/2)^2 + (v - N/2)^2\right].$$

The Laplacian-filtered image in the spatial domain is obtained by computing the inverse Fourier transform of H (u, v) F (u, v):

$$\nabla^2 f(x, y) = \Im^{-1}\left\{-\left[(u - M/2)^2 + (v - N/2)^2\right]F(u, v)\right\}.$$

Conversely, computing the Laplacian in the spatial domain and computing the Fourier transform of the result is equivalent to multiplying F(u, v) by H(u, v). We express this dual relationship in the familiar Fourier-transform-pair notation

$$\nabla^2 f(x, y) \Leftrightarrow -\left[(u - M/2)^2 + (v - N/2)^2\right]F(u, v).$$

The spatial domain Laplacian filter function obtained by taking the inverse Fourier transform of Eq. has some interesting properties, as Fig.7 shows. Figure 7(a) is a 3-D perspective plot. The function is centered at (M/2, N/2), and its value at the top of the dome is zero. All other values are negative. Figure 7(b) shows H (u, v) as an image, also centered. Figure 7(c) is the Laplacian in the spatial domain, obtained by multiplying by H (u, v) by $(-1)^{u+v}$, taking the inverse Fourier transform, and multiplying the real part of the result by $(-1)^{x+y}$. Figure 7(d) is a zoomed section at about the origin of Fig.7(c).' Figure 7(e) is a horizontal gray-level profile passing through the center of the zoomed section. Finally, Fig.7 (f) shows the mask to implement the definition of the discrete Laplacian in the spatial domain.
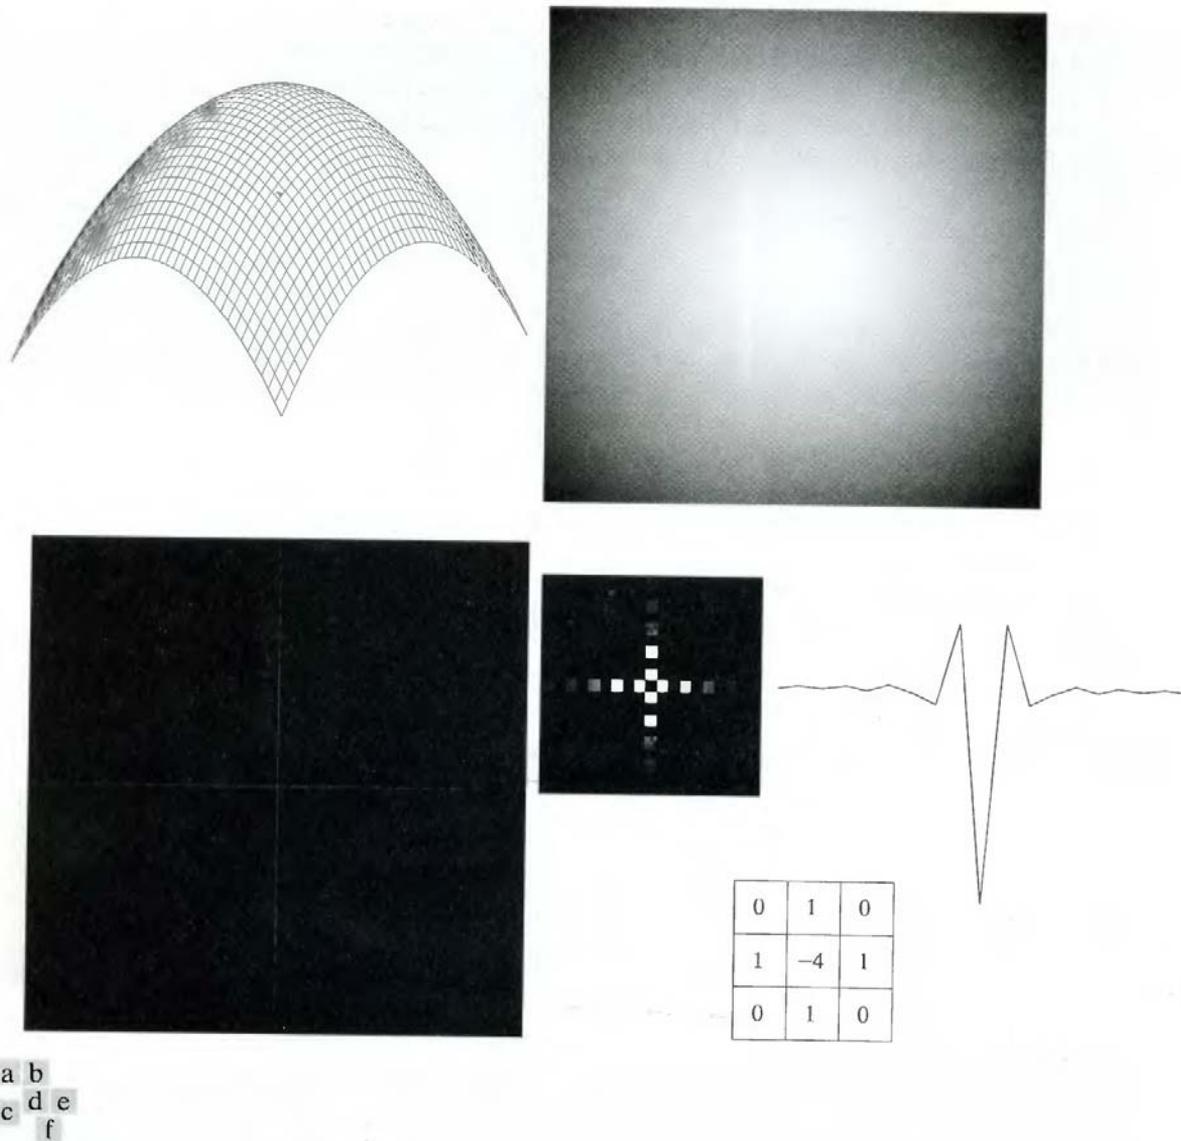
Fig.7 (a) 3-D plot of Laplacian in the frequency domain, (b) Image representation of (a), (c) Laplacian in the spatial domain obtained from the inverse DFT of (b) (d) Zoomed section of the origin of (c). (e) Gray-level profile through the center of (d). (f) Laplacian mask

 A horizontal profile through the center of this mask has the same basic shape as the profile in Fig. 7(e) (that is, a negative value between two smaller positive values). We form an enhanced image g(x, y) by subtracting the Laplacian from the original image:

$$g(x, y) = f(x, y) - \nabla^2 f(x, y).$$

8. Write about high boost and high frequency filtering.

High-Boost Filtering and High-Frequency Emphasis Filtering:

All the filtered images have one thing in common: Their average background intensity has been reduced to near black. This is due to the fact that the highpass filters we applied to those images eliminate the zero-frequency component of their Fourier transforms. In fact, enhancement using the Laplacian does precisely this, by adding back the entire image to the filtered result. Sometimes it is advantageous to increase the contribution made by the original image to the overall filtered result. This approach, called high-boost filtering, is a generalization of unsharp masking. Unsharp masking consists simply of generating a sharp image by subtracting from an image a blurred version of itself. Using frequency domain terminology, this means obtaining a highpass-filtered image by subtracting from the image a lowpass-filtered version of itself. That is

$$f_{hp}'(x, y) = f(x, y) - f_{lp}(x, y).$$

High-boost filtering generalizes this by multiplying f (x, y) by a constant A > 1:

$$f_{hb} = Af(x, y) - f_{lp}(x, y).$$

Thus, high-boost filtering gives us the flexibility to increase the contribution made by the image to the overall enhanced result. This equation may be written as

$$f_{hb}(x, y) = (A - 1)f(x, y) + f(x, y) - f_{lp}(x, y).$$

Then, using above Eq. we obtain

$$f_{hb}(x, y) = (A - 1)f(x, y) + f_{hp}(x, y).$$

This result is based on a highpass rather than a lowpass image. When A = 1, high-boost filtering reduces to regular highpass filtering. As A increases past 1, the contribution made by the image itself becomes more dominant.

We have $F_{hp}(u,v) = F(u,v) - F_{lp}(u,v)$. But $F_{lp}(u,v) = H_{lp}(u,v)F(u,v)$, where H_{lp} is the transfer function of a lowpass filter. Therefore, unsharp masking can be implemented directly in the frequency domain by using the composite filter

$$H_{hp}(u, v) = 1 - H_{lp}(u, v).$$

Similarly, high-boost filtering can be implemented with the composite filter

$$H_{hb}(u, v) = (A - 1) + H_{hp}(u, v)$$

with $A > 1$. The process consists of multiplying this filter by the (centered) transform of the input image and then taking the inverse transform of the product. Multiplication of the real part of this result by $(-1)^{x+y}$ gives us the high-boost filtered image $f_{hb}(x, y)$ in the spatial domain.

9. Explain the concept of homomorphic filtering.

Homomorphic filtering:

The illumination-reflectance model can be used to develop a frequency domain procedure for improving the appearance of an image by simultaneous gray-level range compression and contrast enhancement. An image $f(x, y)$ can be expressed as the product of illumination and reflectance components:

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

Equation above cannot be used directly to operate separately on the frequency components of illumination and reflectance because the Fourier transform of the product of two functions is not separable; in other words,

$$\Im\{f(x, y)\} \neq \Im\{i(x, y)\}\Im\{r(x, y)\}.$$

Suppose, however, that we define

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

Then

$$\Im\{z(x, y)\} = \Im\{\ln f(x, y)\}$$
$$= \Im\{\ln i(x, y)\} + \Im\{\ln r(x, y)\}$$

or

$$Z(u, v) = F_i(u, v) + F_r(u, v)$$

where F_i (u, v) and F_r (u, v) are the Fourier transforms of ln i(x, y) and ln r(x, y), respectively. If we process Z (u, v) by means of a filter function H (u, v) then, from

$$S(u, v) = H(u, v)Z(u, v)$$
$$= H(u, v)F_i(u, v) + H(u, v)F_r(u, v)$$

where S (u, v) is the Fourier transform of the result. In the spatial domain,

$$s(x, y) = \Im^{-1}\{S(u, v)\}$$
$$= \Im^{-1}\{H(u, v)F_i(u, v)\} + \Im^{-1}\{H(u, v)F_r(u, v)\}.$$

By letting

$$i'(x, y) = \Im^{-1}\{H(u, v)F_i(u, v)\}$$

and

$$r'(x, y) = \Im^{-1}\{H(u, v)F_r(u, v)\},$$

Now we have

$$s(x, y) = i'(x, y) + r'(x, y).$$

Finally, as z (x, y) was formed by taking the logarithm of the original image f (x, y), the inverse (exponential) operation yields the desired enhanced image, denoted by g(x, y); that is,

$$g(x, y) = e^{s(x,y)}$$
$$= e^{i'(x,y)} \cdot e^{r'(x,y)}$$
$$= i_0(x, y)r_0(x, y)$$

where

$$i_0(x, y) = e^{i'(x,y)}$$

$$f(x, y) \Rightarrow \boxed{\ln} \Rightarrow \boxed{DFT} \Rightarrow \boxed{H(u, v)} \Rightarrow \boxed{(DFT)^{-1}} \Rightarrow \boxed{\exp} \Rightarrow g(x, y)$$
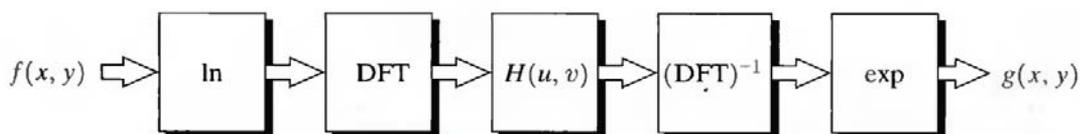
Fig.9.1 Homomorphic filtering approach for image enhancement

and

$$r_0(x, y) = e^{r'(x, y)}$$

are the illumination and reflectance components of the output image. The enhancement approach using the foregoing concepts is summarized in Fig. 9.1. This method is based on a special case of a class of systems known as homomorphic systems. In this particular application, the key to the approach is the separation of the illumination and reflectance components achieved. The homomorphic filter function H (u, v) can then operate on these components separately.

The illumination component of an image generally is characterized by slow spatial variations, while the reflectance component tends to vary abruptly, particularly at the junctions of dissimilar objects. These characteristics lead to associating the low frequencies of the Fourier transform of the logarithm of an image with illumination and the high frequencies with reflectance. Although these associations are rough approximations, they can be used to advantage in image enhancement.

A good deal of control can be gained over the illumination and reflectance components with a homomorphic filter. This control requires specification of a filter function H (u, v) that affects the low- and high-frequency components of the Fourier transform in different ways. Figure 9.2 shows a cross section of such a filter. If the parameters γ_L and γ_H are chosen so that $\gamma_L < 1$ and $\gamma_H > 1$, the filter function shown in Fig. 9.2 tends to decrease the contribution made by the low frequencies (illumination) and amplify the contribution made by high frequencies (reflectance). The net result is simultaneous dynamic range compression and contrast enhancement.
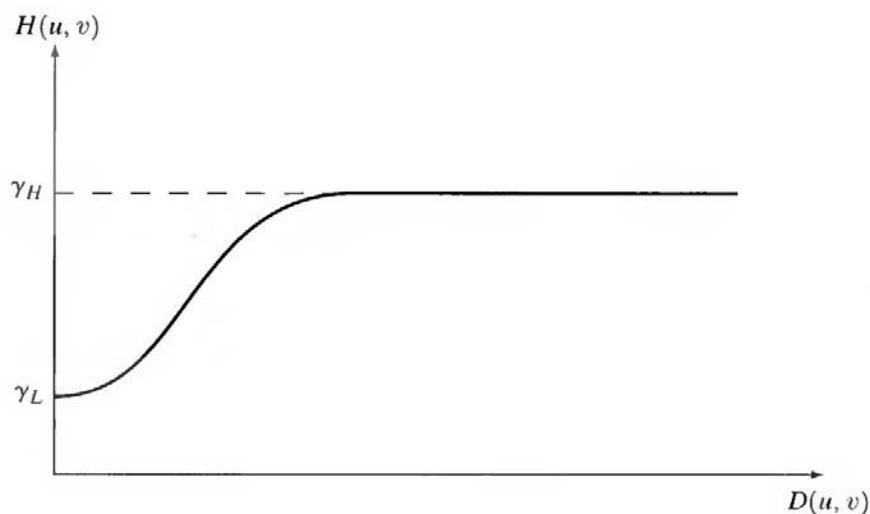
Fig.9.2 Cross section of a circularly symmetric filter function D (u. v) is the distance from the origin of the centered transform.

1. What are the derivative operators useful in image segmentation? Explain their role in segmentation.

Gradient operators:

First-order derivatives of a digital image are based on various approximations of the 2-D gradient. The gradient of an image $f(x, y)$ at location (x, y) is defined as the vector

$$\nabla \mathbf{f} = \begin{bmatrix} G_x \\ G_y \end{bmatrix} = \begin{bmatrix} \dfrac{\partial f}{\partial x} \\ \dfrac{\partial f}{\partial y} \end{bmatrix}.$$

It is well known from vector analysis that the gradient vector points in the direction of maximum rate of change of f at coordinates (x, y). An important quantity in edge detection is the magnitude of this vector, denoted by ∇f, where

$$\nabla f = \mathrm{mag}(\nabla \mathbf{f}) = \left[G_x^2 + G_y^2 \right]^{1/2}.$$

This quantity gives the maximum rate of increase of $f(x, y)$ per unit distance in the direction of ∇f. It is a common (although not strictly correct) practice to refer to ∇f also as the gradient. The direction of the gradient vector also is an important quantity. Let $\alpha(x, y)$ represent the direction angle of the vector ∇f at (x, y). Then, from vector analysis,

$$\alpha(x, y) = \tan^{-1}\left(\frac{G_y}{G_x} \right)$$

where the angle is measured with respect to the x-axis. The direction of an edge at (x, y) is perpendicular to the direction of the gradient vector at that point. Computation of the gradient of an image is based on obtaining the partial derivatives $\partial f/\partial x$ and $\partial f/\partial y$ at every pixel location. Let the 3x3 area shown in Fig. 1.1 (a) represent the gray levels in a neighborhood of an image. One of the simplest ways to implement a first-order partial derivative at point z_5 is to use the following Roberts cross-gradient operators:

$$G_x = (z_9 - z_5)$$

and

$$G_y = (z_8 - z_6).$$

These derivatives can be implemented for an entire image by using the masks shown in Fig. 1.1(b). Masks of size 2 X 2 are awkward to implement because they do not have a clear center. An approach using masks of size 3 X 3 is given by
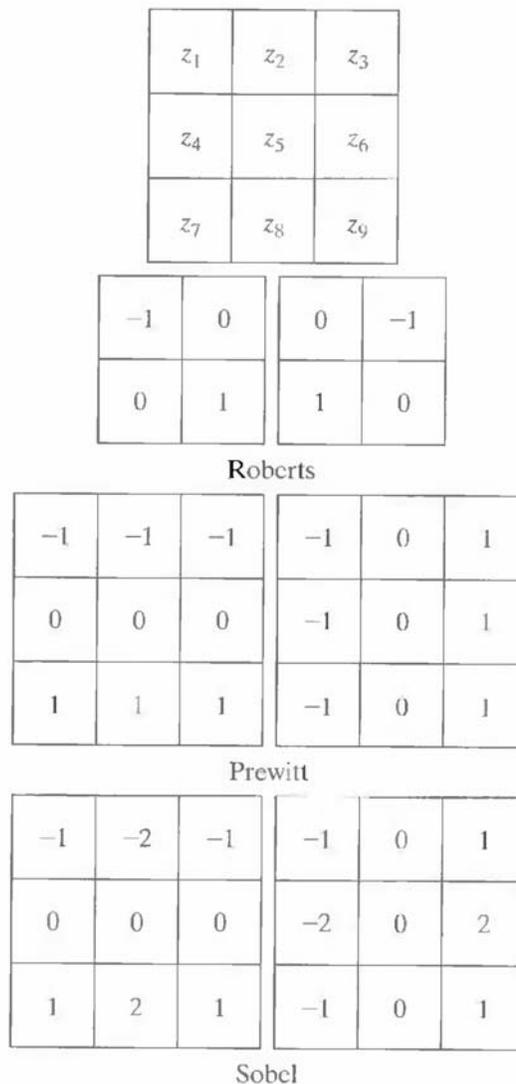
Fig.1.1 A 3 X 3 region of an image (the z's are gray-level values) and various masks used to compute the gradient at point labeled z_5.

$$G_x = (z_7 + z_8 + z_9) - (z_1 + z_2 + z_3)$$

and

$$G_y = (z_3 + z_6 + z_9) - (z_1 + z_4 + z_7).$$

A weight value of 2 is used to achieve some smoothing by giving more importance to the center point. Figures 1.1(f) and (g), called the Sobel operators, and are used to implement these two equations. The Prewitt and Sobel operators are among the most used in practice for computing digital gradients. The Prewitt masks are simpler to implement than the Sobel masks, but the latter have slightly superior noise-suppression characteristics, an important issue when dealing with derivatives. Note that the coefficients in all the masks shown in Fig. 1.1 sum to 0, indicating that they give a response of 0 in areas of constant gray level, as expected of a derivative operator.

The masks just discussed are used to obtain the gradient components G_x and G_y. Computation of the gradient requires that these two components be combined. However, this implementation is not always desirable because of the computational burden required by squares and square roots. An approach used frequently is to approximate the gradient by absolute values:

$$\nabla f \approx |G_x| + |G_y|.$$

This equation is much more attractive computationally, and it still preserves relative changes in gray levels. However, this is not an issue when masks such as the Prewitt and Sobel masks are used to compute Gx and Gy.

It is possible to modify the 3 X 3 masks in Fig. 1.1 so that they have their strongest responses along the diagonal directions. The two additional Prewitt and Sobel masks for detecting discontinuities in the diagonal directions are shown in Fig. 1.2.

Fig.1.2 Prewitt and Sobel masks for detecting diagonal edges

The Laplacian:

The Laplacian of a 2-D function f(x, y) is a second-order derivative defined as

$$\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}.$$

For a 3 X 3 region, one of the two forms encountered most frequently in practice is

$$\nabla^2 f = 4z_5 - \left(z_2 + z_4 + z_6 + z_8\right)$$

Fig.1.3 Laplacian masks used to implement Eqns. above.

where the z's are defined in Fig. 1.1(a). A digital approximation including the diagonal neighbors is given by

$$\nabla^2 f = 8z_5 - (z_1 + z_2 + z_3 + z_4 + z_6 + z_7 + z_8 + z_9).$$

Masks for implementing these two equations are shown in Fig. 1.3. We note from these masks that the implementations of Eqns. are isotropic for rotation increments of 90° and 45°, respectively.

2. Write about edge detection.

Intuitively, an edge is a set of connected pixels that lie on the boundary between two regions. Fundamentally, an edge is a "local" concept whereas a region boundary, owing to the way it is defined, is a more global idea. A reasonable definition of "edge" requires the ability to measure gray-level transitions in a meaningful way. We start by modeling an edge intuitively. This will lead us to formalism in which "meaningful" transitions in gray levels can be measured. Intuitively, an ideal edge has the properties of the model shown in Fig. 2.1(a). An ideal edge according to this model is a set of connected pixels (in the vertical direction here), each of which is located at an orthogonal step transition in gray level (as shown by the horizontal profile in the figure).

In practice, optics, sampling, and other image acquisition imperfections yield edges that are blurred, with the degree of blurring being determined by factors such as the quality of the image acquisition system, the sampling rate, and illumination conditions under which the image is acquired. As a result, edges are more closely modeled as having a "ramp like" profile, such as the one shown in Fig.2.1 (b).
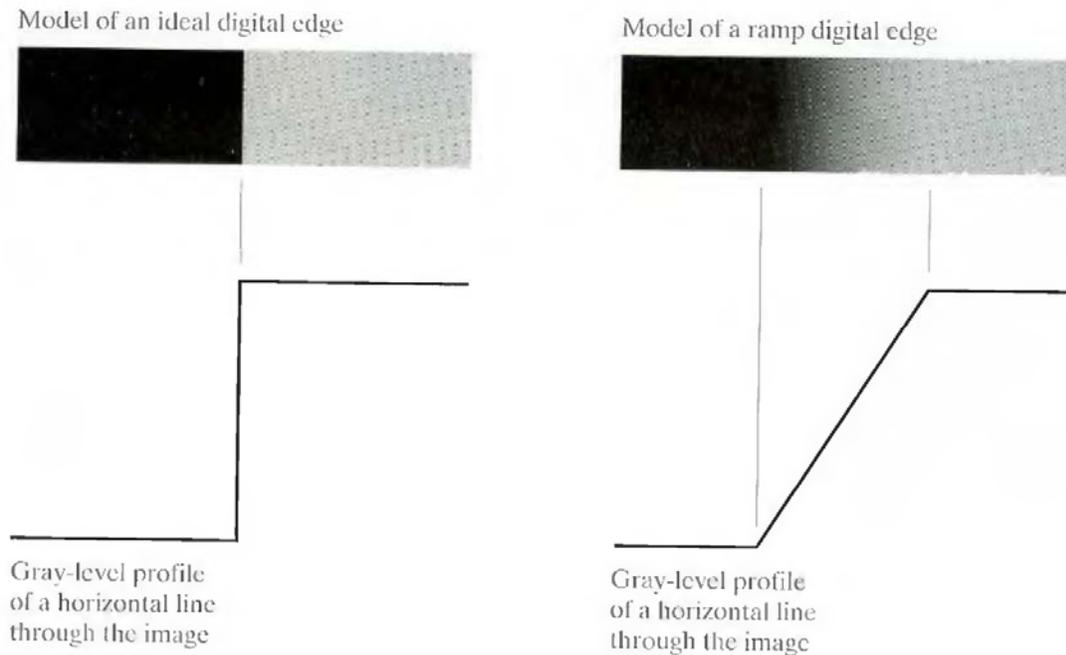
Fig.2.1 (a) Model of an ideal digital edge (b) Model of a ramp edge. The slope of the ramp is proportional to the degree of blurring in the edge.

The slope of the ramp is inversely proportional to the degree of blurring in the edge. In this model, we no longer have a thin (one pixel thick) path. Instead, an edge point now is any point contained in the ramp, and an edge would then be a set of such points that are connected. The "thickness" of the edge is determined by the length of the ramp, as it transitions from an initial to a final gray level. This length is determined by the slope, which, in turn, is determined by the degree of blurring. This makes sense: Blurred edges lend to be thick and sharp edges tend to be thin. Figure 2.2(a) shows the image from which the close-up in Fig. 2.1(b) was extracted. Figure 2.2(b) shows a horizontal gray-level profile of the edge between the two regions. This figure also shows the first and second derivatives of the gray-level profile. The first derivative is positive at the points of transition into and out of the ramp as we move from left to right along the profile; it is constant for points in the ramp; and is zero in areas of constant gray level. The second derivative is positive at the transition associated with the dark side of the edge, negative at the transition associated with the light side of the edge, and zero along the ramp and in areas of constant gray level. The signs of the derivatives in Fig. 2.2(b) would be reversed for an edge that transitions from light to dark.

We conclude from these observations that the magnitude of the first derivative can be used to detect the presence of an edge at a point in an image (i.e. to determine if a point is on a ramp). Similarly, the sign of the second derivative can be used to determine whether an edge pixel lies

on the dark or light side of an edge. We note two additional properties of the second derivative around an edge: A) It produces two values for every edge in an image (an undesirable feature); and B) an imaginary straight line joining the extreme positive and negative values of the second derivative would cross zero near the midpoint of the edge. This zero-crossing property of the second derivative is quite useful for locating the centers of thick edges.
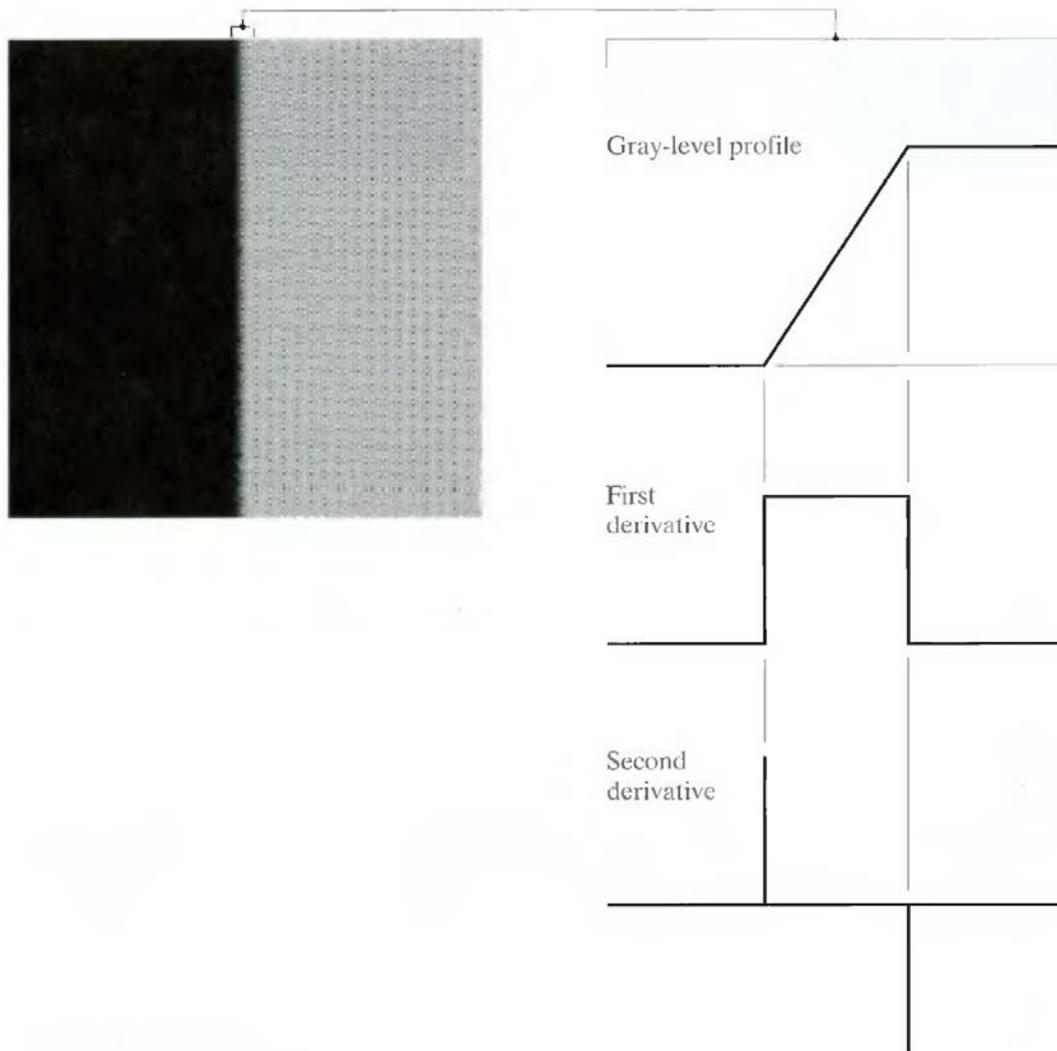
Fig.2.2 (a) Two regions separated by a vertical edge (b) Detail near the edge, showing a gray-level profile, and the first and second derivatives of the profile.

3. Explain about the edge linking procedures.

The different methods for edge linking are as follows

(i) Local processing

(ii) Global processing via the Hough Transform

(iii) Global processing via graph-theoretic techniques.

(i) Local Processing:

One of the simplest approaches for linking edge points is to analyze the characteristics of pixels in a small neighborhood (say, 3 X 3 or 5 X 5) about every point (x, y) in an image that has been labeled an edge point. All points that are similar according to a set of predefined criteria are linked, forming an edge of pixels that share those criteria.

The two principal properties used for establishing similarity of edge pixels in this kind of analysis are (1) the strength of the response of the gradient operator used to produce the edge pixel; and (2) the direction of the gradient vector. The first property is given by the value of ∇f.

Thus an edge pixel with coordinates (x_o, y_o) in a predefined neighborhood of (x, y), is similar in magnitude to the pixel at (x, y) if

$$\left| \nabla f(x, y) - \nabla f(x_0, y_0) \right| \le E$$

The direction (angle) of the gradient vector is given by Eq. An edge pixel at (x_o, y_o) in the predefined neighborhood of (x, y) has an angle similar to the pixel at (x, y) if

$$\left| \alpha(x, y) - \alpha(x_0, y_0) \right| < A$$

where A is a nonnegative angle threshold. The direction of the edge at (x, y) is perpendicular to the direction of the gradient vector at that point.

A point in the predefined neighborhood of (x, y) is linked to the pixel at (x, y) if both magnitude and direction criteria are satisfied. This process is repeated at every location in the image. A record must be kept of linked points as the center of the neighborhood is moved from pixel to pixel. A simple bookkeeping procedure is to assign a different gray level to each set of linked edge pixels.

(ii) Global processing via the Hough Transform:

In this process, points are linked by determining first if they lie on a curve of specified shape. We now consider global relationships between pixels. Given n points in an image, suppose that we want to find subsets of these points that lie on straight lines. One possible solution is to first find all lines determined by every pair of points and then find all subsets of points that are close to particular lines. The problem with this procedure is that it involves finding $n(n - 1)/2 \sim n^2$ lines and then performing $(n)(n(n - 1))/2 \sim n^3$ comparisons of every point to all lines. This approach is computationally prohibitive in all but the most trivial applications.

Hough [1962] proposed an alternative approach, commonly referred to as the Hough transform. Consider a point (x_i, y_i) and the general equation of a straight line in slope-intercept form, $y_i = a.x_i + b$. Infinitely many lines pass through (x_i, y_i) but they all satisfy the equation $y_i = a.x_i + b$ for varying values of a and b. However, writing this equation as $b = -a.x_i + y_i$, and considering the ab-plane (also called parameter space) yields the equation of a single line for a fixed pair (x_i, y_i). Furthermore, a second point (x_j, y_j) also has a line in parameter space associated with it, and this line intersects the line associated with (x_i, y_i) at (a', b'), where a' is the slope and b' the intercept of the line containing both (x_i, y_i) and (x_j, y_j) in the xy-plane. In fact, all points contained on this line have lines in parameter space that intersect at (a', b'). Figure 3.1 illustrates these concepts.
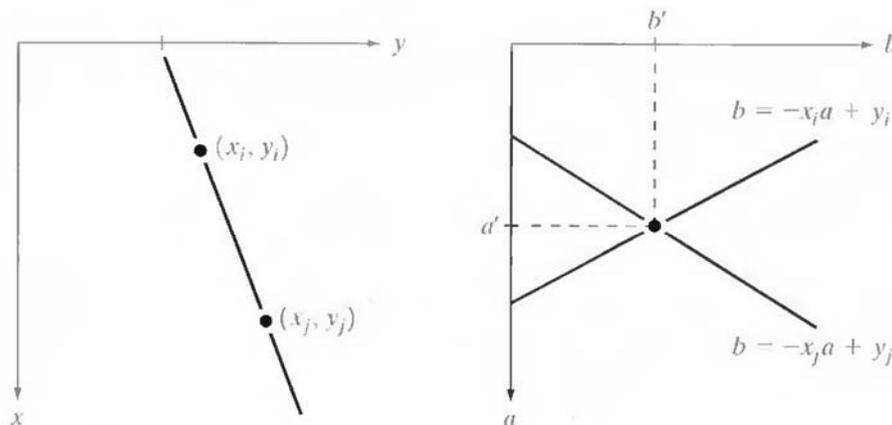
Fig.3.1 (a) xy-plane (b) Parameter space
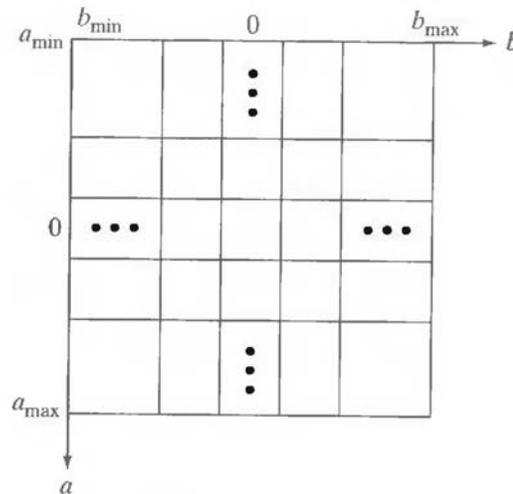
Fig.3.2 Subdivision of the parameter plane for use in the Hough transform

The computational attractiveness of the Hough transform arises from subdividing the parameter space into so-called accumulator cells, as illustrated in Fig. 3.2, where (a_{max} , a_{min}) and (b_{max} , b_{min}), are the expected ranges of slope and intercept values. The cell at coordinates (i, j), with accumulator value A(i, j), corresponds to the square associated with parameter space coordinates (a_i , b_i).

Initially, these cells are set to zero. Then, for every point (x_k, y_k) in the image plane, we let the parameter a equal each of the allowed subdivision values on the fl-axis and solve for the corresponding b using the equation b = - x_k a + y_k .The resulting b's are then rounded off to the nearest allowed value in the b-axis. If a choice of a_p results in solution b_q, we let A (p, q) = A (p, q) + 1. At the end of this procedure, a value of Q in A (i, j) corresponds to Q points in the xy-plane lying on the line y = a_i x + b_j. The number of subdivisions in the ab-plane determines the accuracy of the co linearity of these points. Note that subdividing the a axis into K increments gives, for every point (x_k, y_k), K values of b corresponding to the K possible values of a. With n image points, this method involves nK computations. Thus the procedure just discussed is linear in n, and the product nK does not approach the number of computations discussed at the beginning unless K approaches or exceeds n.

　　　　　A problem with using the equation y = ax + b to represent a line is that the slope approaches infinity as the line approaches the vertical. One way around this difficulty is to use the normal representation of a line:

$$x \cos\theta + y \sin\theta = \rho$$

Figure 3.3(a) illustrates the geometrical interpretation of the parameters used. The use of this representation in constructing a table of accumulators is identical to the method discussed for the slope-intercept representation. Instead of straight lines, however, the loci are sinusoidal curves in the $\rho\theta$ -plane. As before, Q collinear points lying on a line $x \cos\theta_j + y \sin\theta_j = \rho$, yield Q sinusoidal curves that intersect at (ρ_i, θ_j) in the parameter space. Incrementing θ and solving for the corresponding p gives Q entries in accumulator A (i, j) associated with the cell determined by (ρ_i, θ_j). Figure 3.3 (b) illustrates the subdivision of the parameter space.
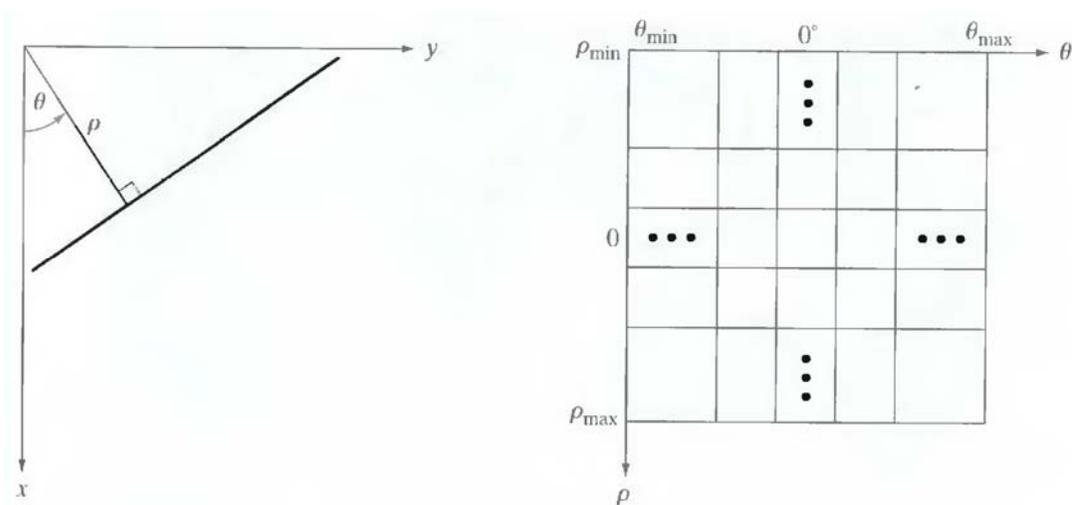
Fig.3.3 (a) Normal representation of a line (b) Subdivision of the $\rho\theta$-plane into cells

The range of angle θ is $\pm 90°$, measured with respect to the x-axis. Thus with reference to Fig. 3.3 (a), a horizontal line has $\theta = 0°$, with ρ being equal to the positive x-intercept. Similarly, a vertical line has $\theta = 90°$, with p being equal to the positive y-intercept, or $\theta = - 90°$, with ρ being equal to the negative y-intercept.

(iii) Global processing via graph-theoretic techniques

In this process we have a global approach for edge detection and linking based on representing edge segments in the form of a graph and searching the graph for low-cost paths that correspond to significant edges. This representation provides a rugged approach that performs well in the presence of noise.
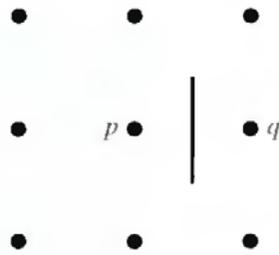
Fig.3.4 Edge clement between pixels p and q

We begin the development with some basic definitions. A graph G = (N,U) is a finite, nonempty set of nodes N, together with a set U of unordered pairs of distinct elements of N. Each pair (n_i, n_j) of U is called an arc. A graph in which the arcs are directed is called a directed graph. If an arc is directed from node n_i to node n_j, then n_j is said to be a successor of the parent node n_i. The process of identifying the successors of a node is called expansion of the node. In each graph we define levels, such that level 0 consists of a single node, called the start or root node, and the nodes in the last level are called goal nodes. A cost c (n_i, n_j) can be associated with every arc (n_i, n_j). A sequence of nodes n_1, n_2... n_k, with each node n_i being a successor of node n_{i-1} is called a path from n_1 to n_k. The cost of the entire path is

$$c = \sum_{i=2}^{k} c(n_{i-1}, n_i).$$

The following discussion is simplified if we define an edge element as the boundary between two pixels p and q, such that p and q are 4-neighbors, as Fig.3.4 illustrates. Edge elements are identified by the xy-coordinates of points p and q. In other words, the edge element in Fig. 3.4 is defined by the pairs (x_p, y_p) (x_q, y_q). Consistent with the definition an edge is a sequence of connected edge elements.

We can illustrate how the concepts just discussed apply to edge detection using the 3 X 3 image shown in Fig. 3.5 (a). The outer numbers are pixel

a b c
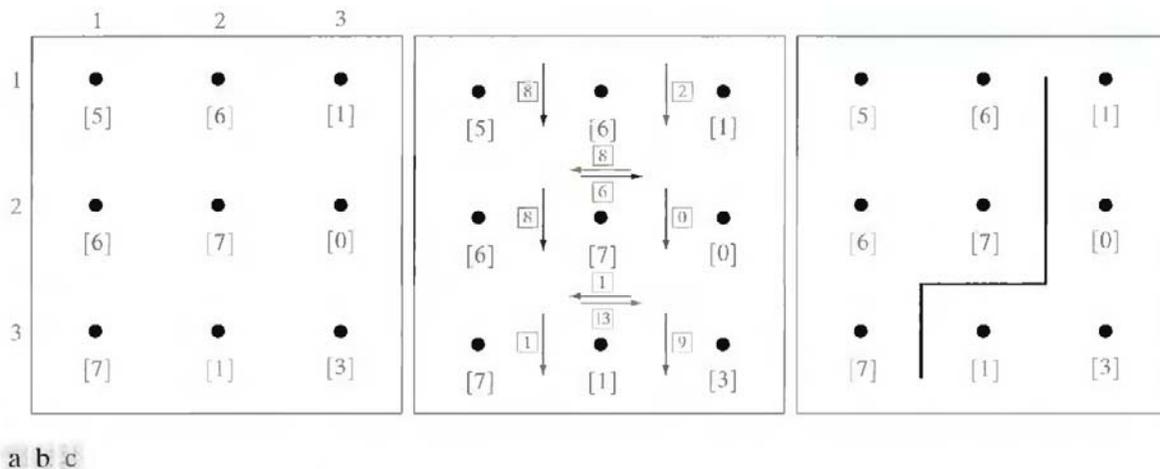
Fig.3.5 (a) A 3 X 3 image region, (b) Edge segments and their costs, (c) Edge corresponding to the lowest-cost path in the graph shown in Fig. 3.6

coordinates and the numbers in brackets represent gray-level values. Each edge element, defined by pixels p and q, has an associated cost, defined as

$$c(p, q) = H - \left[f(p) - f(q) \right]$$

where H is the highest gray-level value in the image (7 in this case), and f(p) and f(q) are the gray-level values of p and q, respectively. By convention, the point p is on the right-hand side of the direction of travel along edge elements. For example, the edge segment (1, 2) (2, 2) is between points (1, 2) and (2, 2) in Fig. 3.5 (b). If the direction of travel is to the right, then p is the point with coordinates (2, 2) and q is point with coordinates (1, 2); therefore, c (p, q) = 7 - [7 - 6] = 6. This cost is shown in the box below the edge segment. If, on the other hand, we are traveling to the left between the same two points, then p is point (1, 2) and q is (2, 2). In this case the cost is 8, as shown above the edge segment in Fig. 3.5(b). To simplify the discussion, we assume that edges start in the top row and terminate in the last row, so that the first element of an edge can be only between points (1, 1), (1, 2) or (1, 2), (1, 3). Similarly, the last edge element has to be between points (3, 1), (3, 2) or (3, 2), (3, 3). Keep in mind that p and q are 4-neighbors, as noted earlier. Figure 3.6 shows the graph for this problem. Each node (rectangle) in the graph corresponds to an edge element from Fig. 3.5. An arc exists between two nodes if the two corresponding edge elements taken in succession can be part of an edge.
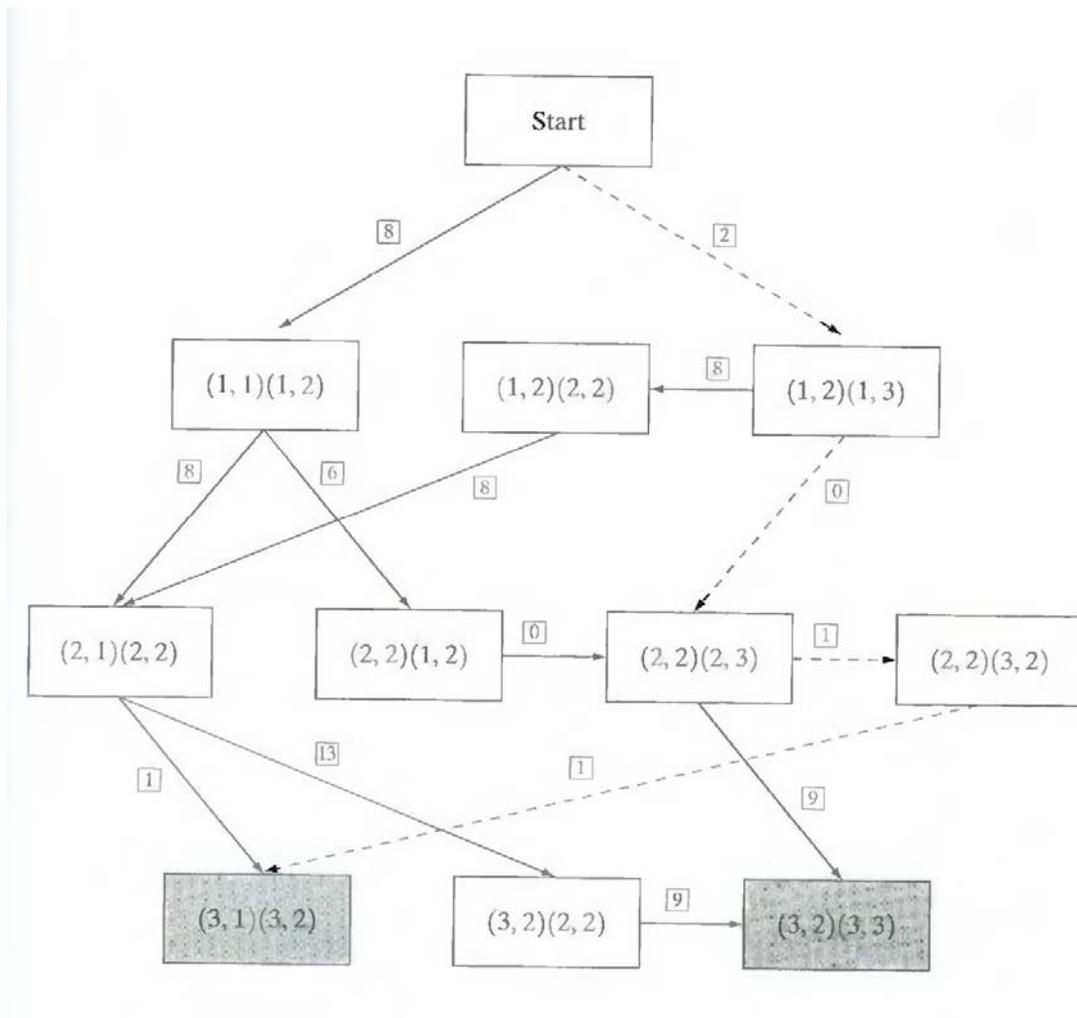
Fig. 3.6 Graph for the image in Fig.3.5 (a). The lowest-cost path is shown dashed.

As in Fig. 3.5 (b), the cost of each edge segment, is shown in a box on the side of the arc leading into the corresponding node. Goal nodes are shown shaded. The minimum cost path is shown dashed, and the edge corresponding to this path is shown in Fig. 3.5 (c).

Digital Image Processing   **Question & Answers**

4. What is thresholding? Explain about global thresholding.

Thresholding:

Because of its intuitive properties and simplicity of implementation, image thresholding enjoys a central position in applications of image segmentation.

Global Thresholding:

The simplest of all thresholding techniques is to partition the image histogram by using a single global threshold, T. Segmentation is then accomplished by scanning the image pixel by pixel and labeling each pixel as object or back-ground, depending on whether the gray level of that pixel is greater or less than the value of T. As indicated earlier, the success of this method depends entirely on how well the histogram can be partitioned.
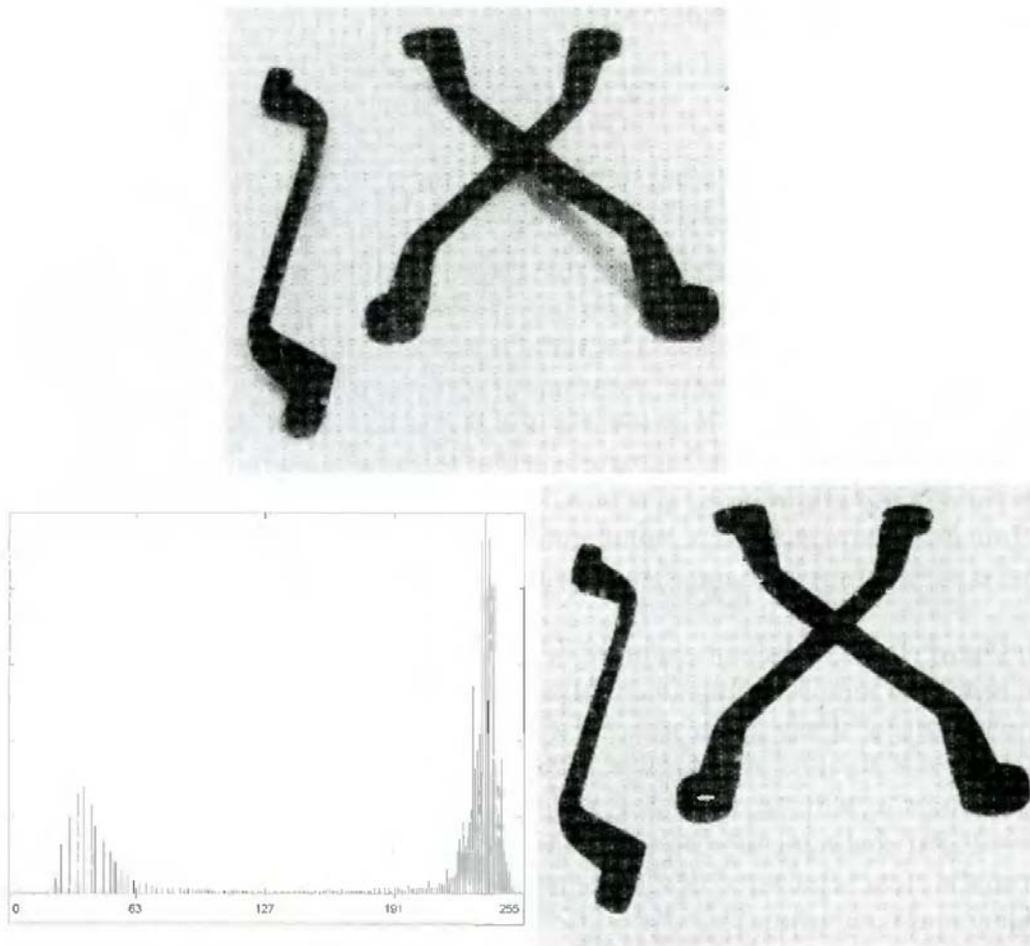
a
b c **Fig.4.1 FIGURE 10.28 (a) Original image, (b) Image histogram, (c) Result of global thresholding with T midway between the maximum and minimum gray levels.**

Figure 4.1(a) shows a simple image, and Fig. 4.1(b) shows its histogram. Figure 4.1(c) shows the result of segmenting Fig. 4.1(a) by using a threshold T midway between the maximum and minimum gray levels. This threshold achieved a "clean" segmentation by eliminating the shadows and leaving only the objects themselves. The objects of interest in this case are darker than the background, so any pixel with a gray level $\leq$ T was labeled black (0), and any pixel with a gray level $\geq$ T was labeled white (255).The key objective is merely to generate a binary image, so the black-white relationship could be reversed. The type of global thresholding just described can be expected to be successful in highly controlled environments. One of the areas in which this often is possible is in industrial inspection applications, where control of the illumination usually is feasible.

The threshold in the preceding example was specified by using a heuristic approach, based on visual inspection of the histogram. The following algorithm can be used to obtain T automatically:

1. Select an initial estimate for T.

2. Segment the image using T. This will produce two groups of pixels: G_1 consisting of all pixels with gray level values >T and G_2 consisting of pixels with values < T.

3. Compute the average gray level values μ_1 and μ_2 for the pixels in regions G_1 and G_2.

4. Compute a new threshold value:

$$T = \frac{1}{2}(\mu_1 + \mu_2).$$

5. Repeat steps 2 through 4 until the difference in T in successive iterations is smaller than a predefined parameter T_o.

When there is reason to believe that the background and object occupy comparable areas in the image, a good initial value for T is the average gray level of the image. When objects are small compared to the area occupied by the background (or vice versa), then one group of pixels will dominate the histogram and the average gray level is not as good an initial choice. A more appropriate initial value for T in cases such as this is a value midway between the maximum and minimum gray levels. The parameter T_o is used to stop the algorithm after changes become small in terms of this parameter. This is used when speed of iteration is an important issue.

5. Explain about basic adaptive thresholding process used in image segmentation.

Basic Adaptive Thresholding:

Imaging factors such as uneven illumination can transform a perfectly segmentable histogram into a histogram that cannot be partitioned effectively by a single global threshold. An approach for handling such a situation is to divide the original image into subimages and then utilize a

different threshold to segment each subimage. The key issues in this approach are how to subdivide the image and how to estimate the threshold for each resulting subimage. Since the threshold used for each pixel depends on the location of the pixel in terms of the subimages, this type of thresholding is adaptive.
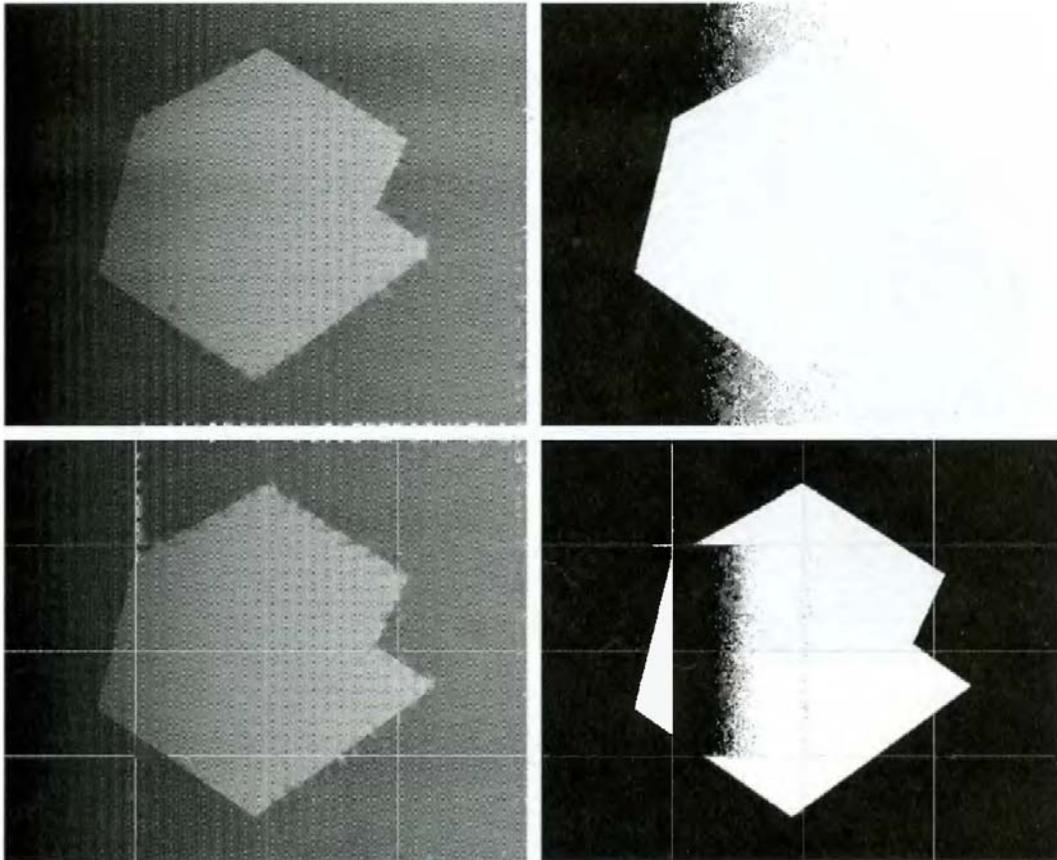
Fig.5 (a) Original image, (b) Result of global thresholding. (c) Image subdivided into individual subimages (d) Result of adaptive thresholding.

We illustrate adaptive thresholding with a example. Figure 5(a) shows the image, which we concluded could not be thresholded effectively with a single global threshold. In fact, Fig. 5(b) shows the result of thresholding the image with a global threshold manually placed in the valley of its histogram. One approach to reduce the effect of nonuniform illumination is to subdivide the image into smaller subimages, such that the illumination of each subimage is approximately uniform. Figure 5(c) shows such a partition, obtained by subdividing the image into four equal parts, and then subdividing each part by four again. All the subimages that did not contain a boundary between object and back-ground had variances of less than 75. All subimages containing boundaries had variances in excess of 100. Each subimage with variance greater than 100 was segmented with a threshold computed for that subimage using the algorithm. The initial

value for T in each case was selected as the point midway between the minimum and maximum gray levels in the subimage. All subimages with variance less than 100 were treated as one composite image, which was segmented using a single threshold estimated using the same algorithm. The result of segmentation using this procedure is shown in Fig. 5(d).

With the exception of two subimages, the improvement over Fig. 5(b) is evident. The boundary between object and background in each of the improperly segmented subimages was small and dark, and the resulting histogram was almost unimodal.

6. Explain in detail the threshold selection based on boundary characteristics.

Use of Boundary Characteristics for Histogram Improvement and Local Thresholding:

 It is intuitively evident that the chances of selecting a "good" threshold are enhanced considerably if the histogram peaks are tall, narrow, symmetric, and separated by deep valleys. One approach for improving the shape of histograms is to consider only those pixels that lie on or near the edges between objects and the background. An immediate and obvious improvement is that histograms would be less dependent on the relative sizes of objects and the background. For instance, the histogram of an image composed of a small object on a large background area (or vice versa) would be dominated by a large peak because of the high concentration of one type of pixels.

        If only the pixels on or near the edge between object and the background were used, the resulting histogram would have peaks of approximately the same height. In addition, the probability that any of those given pixels lies on an object would be approximately equal to the probability that it lies on the back-ground, thus improving the symmetry of the histogram peaks.

                Finally, as indicated in the following paragraph, using pixels that satisfy some simple measures based on gradient and Laplacian operators has a tendency to deepen the valley between histogram peaks.

The principal problem with the approach just discussed is the implicit assumption that the edges between objects and background arc known. This information clearly is not available during segmentation, as finding a division between objects and background is precisely what segmentation is all about. However, an indication of whether a pixel is on an edge may be obtained by computing its gradient. In addition, use of the Laplacian can yield information regarding whether a given pixel lies on the dark or light side of an edge. The average value of the Laplacian is 0 at the transition of an edge, so in practice the valleys of histograms formed from

the pixels selected by a gradient/Laplacian criterion can be expected to be sparsely populated. This property produces the highly desirable deep valleys.

The gradient ∇f at any point (x, y) in an image can be found. Similarly, the Laplacian $\nabla^2 f$ can also be found. These two quantities may be used to form a three-level image, as follows:

$$s(x, y) = \begin{cases} 0 & \text{if } \nabla f < T \\ + & \text{if } \nabla f \geq T \quad \text{and} \quad \nabla^2 f \geq 0 \\ - & \text{if } \nabla f \geq T \quad \text{and} \quad \nabla^2 f < 0 \end{cases}$$

where the symbols 0, +, and - represent any three distinct gray levels, T is a threshold, and the gradient and Laplacian are computed at every point (x, y). For a dark object on a light background, the use of the Eqn. produces an image s(x, y) in which (1) all pixels that are not on an edge (as determined by ∇f being less than T) are labeled 0; (2) all pixels on the dark side of an edge are labeled +; and (3) all pixels on the light side of an edge are labeled -. The symbols + and - in Eq. above are reversed for a light object on a dark background. Figure 6.1 shows the labeling produced by Eq. for an image of a dark, underlined stroke written on a light background.

The information obtained with this procedure can be used to generate a segmented, binary image in which l's correspond to objects of interest and 0's correspond to the background. The transition (along a horizontal or vertical scan line) from a light background to a dark object must be characterized by the occurrence of a - followed by a + in s (x, y). The interior of the object is composed of pixels that are labeled either 0 or +. Finally, the transition from the object back to the background is characterized by the occurrence of a + followed by a -. Thus a horizontal or vertical scan line containing a section of an object has the following structure:

$$(\cdots)(-, +)(0 \text{ or } +)(+, -)(\cdots)$$
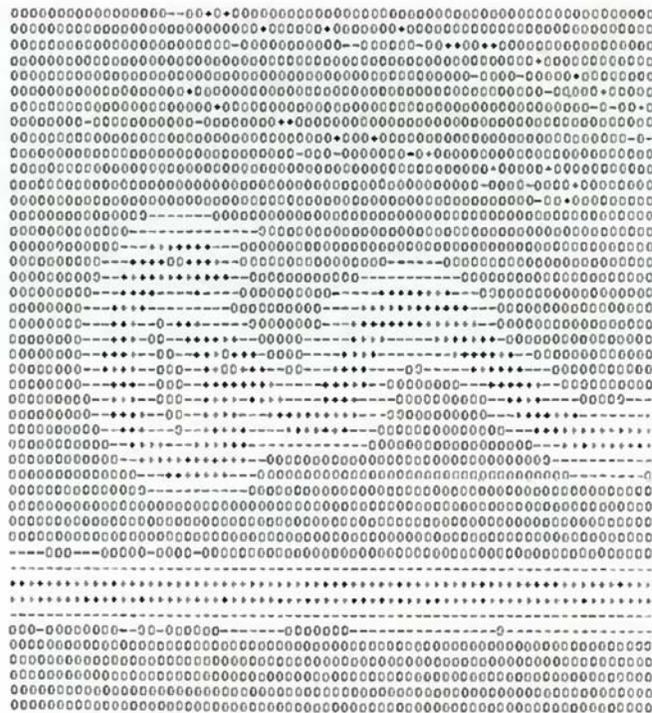
Fig.6.1 Image of a handwritten stroke coded by using Eq. discussed above

where (…) represents any combination of +, -, and 0. The innermost parentheses contain object points and are labeled 1. All other pixels along the same scan line are labeled 0, with the exception of any other sequence of (- or +) bounded by (-, +) and (+, -).
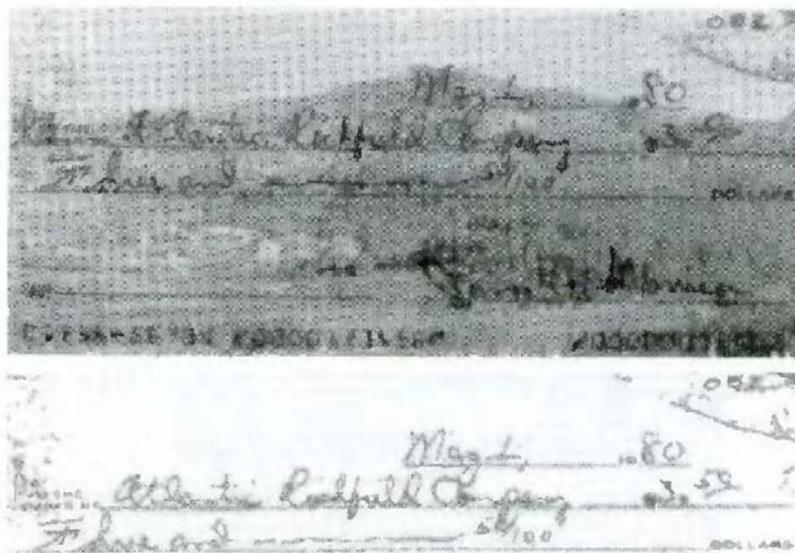
Fig.6.2 (a) Original image, (b) Image segmented by local thresholding.

Figure 6.2 (a) shows an image of an ordinary scenic bank check. Figure 6.3 shows the histogram as a function of gradient values for pixels with gradients greater than 5. Note that this histogram has two dominant modes that are symmetric, nearly of the same height, and arc separated by a distinct valley. Finally, Fig. 6.2(b) shows the segmented image obtained by with T at or near the midpoint of the valley. Note that this example is an illustration of local thresholding, because the value of T was determined from a histogram of the gradient and Laplacian, which are local properties.
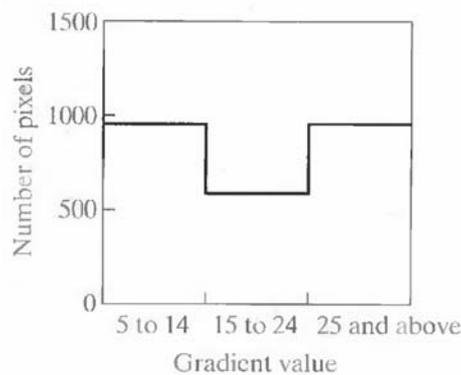
Fig.6.3 Histogram of pixels with gradients greater than 5

7. Explain about region based segmentation.

Region-Based Segmentation:

The objective of segmentation is to partition an image into regions. We approached this problem by finding boundaries between regions based on discontinuities in gray levels, whereas segmentation was accomplished via thresholds based on the distribution of pixel properties, such as gray-level values or color.

Basic Formulation:

Let R represent the entire image region. We may view segmentation as a process that partitions R into n subregions, R_1, R_2..., R_n, such that

(a) $\bigcup\limits_{i=1}^{n} R_i = R.$

(b) R_i is a connected region, $i = 1, 2, \ldots, n.$

(c) $R_i \cap R_j = \emptyset$ for all i and $j, i \neq j.$

(d) $P(R_i) = \text{TRUE for } i = 1, 2, \ldots, n.$

(e) $P(R_i \cup R_j) = \text{FALSE for } i \neq j.$

Here, P (R_i) is a logical predicate defined over the points in set R_i and $\emptyset$ is the null set. Condition (a) indicates that the segmentation must be complete; that is, every pixel must be in a region. Condition (b) requires that points in a region must be connected in some predefined sense. Condition (c) indicates that the regions must be disjoint. Condition (d) deals with the properties that must be satisfied by the pixels in a segmented region—for example P (R_i) = TRUE if all pixels in R_i, have the same gray level. Finally, condition (c) indicates that regions R_i and R_j are different in the sense of predicate P.

Region Growing:

As its name implies, region growing is a procedure that groups pixels or subregions into larger regions based on predefined criteria. The basic approach is to start with a set of "seed" points and from these grow regions by appending to each seed those neighboring pixels that have properties similar to the seed (such as specific ranges of gray level or color). When a priori information is not available, the procedure is to compute at every pixel the same set of properties that ultimately will be used to assign pixels to regions during the growing process. If the result of these computations shows clusters of values, the pixels whose properties place them near the centroid of these clusters can be used as seeds.

The selection of similarity criteria depends not only on the problem under consideration, but also on the type of image data available. For example, the analysis of land-use satellite imagery depends heavily on the use of color. This problem would be significantly more difficult, or even impossible, to handle without the inherent information available in color images. When the images are monochrome, region analysis must be carried out with a set of descriptors based on gray levels and spatial properties (such as moments or texture).

Basically, growing a region should stop when no more pixels satisfy the criteria for inclusion in that region. Criteria such as gray level, texture, and color, are local in nature and do not take into account the "history" of region growth. Additional criteria that increase the power of a region-growing algorithm utilize the concept of size, likeness between a candidate pixel and the pixels grown so far (such as a comparison of the gray level of a candidate and the average gray level of

the grown region), and the shape of the region being grown. The use of these types of descriptors is based on the assumption that a model of expected results is at least partially available.

Figure 7.1 (a) shows an X-ray image of a weld (the horizontal dark region) containing several cracks and porosities (the bright, white streaks running horizontally through the middle of the image). We wish to use region growing to segment the regions of the weld failures. These segmented features could be used for inspection, for inclusion in a database of historical studies, for controlling an automated welding system, and for other numerous applications.
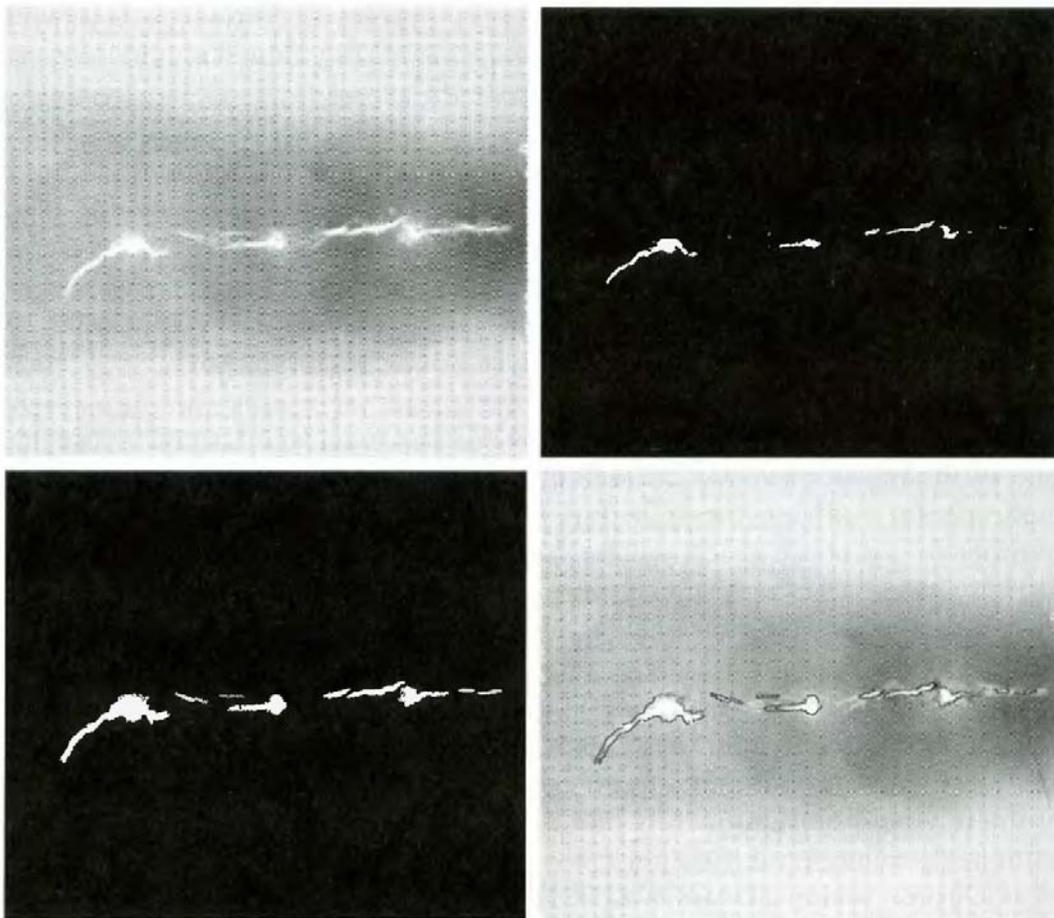
Fig.7.1 (a) Image showing defective welds, (b) Seed points, (c) Result of region growing, (d) Boundaries of segmented ; defective welds (in black).

The first order of business is to determine the initial seed points. In this application, it is known that pixels of defective welds tend to have the maximum allowable digital value B55 in this case). Based on this information, we selected as starting points all pixels having values of 255. The points thus extracted from the original image are shown in Fig. 10.40(b). Note that many of the points are clustered into seed regions.

The next step is to choose criteria for region growing. In this particular example we chose two criteria for a pixel to be annexed to a region: (1) The absolute gray-level difference between any pixel and the seed had to be less than 65. This number is based on the histogram shown in Fig. 7.2 and represents the difference between 255 and the location of the first major valley to the left, which is representative of the highest gray level value in the dark weld region. (2) To be included in one of the regions, the pixel had to be 8-connected to at least one pixel in that region.

If a pixel was found to be connected to more than one region, the regions were merged. Figure 7.1 (c) shows the regions that resulted by starting with the seeds in Fig. 7.2 (b) and utilizing the criteria defined in the previous paragraph. Superimposing the boundaries of these regions on the original image [Fig. 7.1(d)] reveals that the region-growing procedure did indeed segment the defective welds with an acceptable degree of accuracy. It is of interest to note that it was not necessary to specify any stopping rules in this case because the criteria for region growing were sufficient to isolate the features of interest.
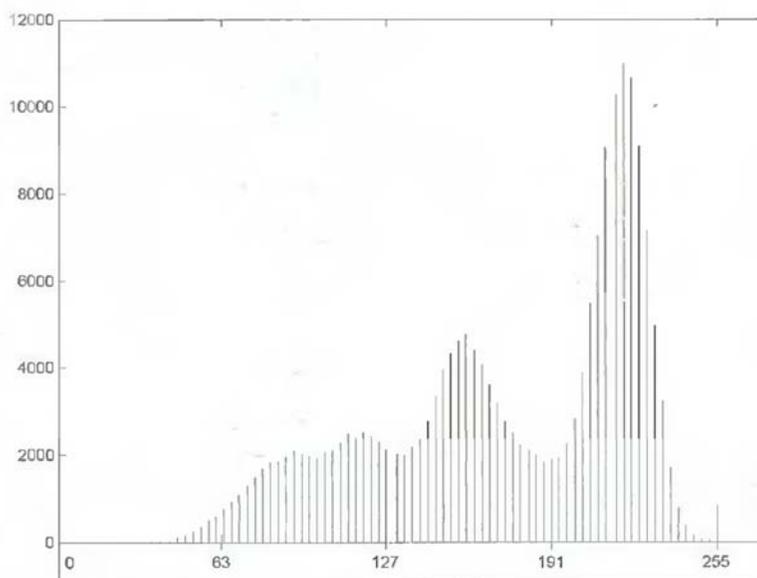
Fig.7.2 Histogram of Fig. 7.1 (a)

Region Splitting and Merging:

The procedure just discussed grows regions from a set of seed points. An alternative is to subdivide an image initially into a set of arbitrary, disjointed regions and then merge and/or split the regions in an attempt to satisfy the conditions. A split and merge algorithm that iteratively works toward satisfying these constraints is developed.

Let R represent the entire image region and select a predicate P. One approach for segmenting R is to subdivide it successively into smaller and smaller quadrant regions so that, for any region R_i, $P(R_i)$ = TRUE. We start with the entire region. If P(R) = FALSE, we divide the image into quadrants. If P is FALSE for any quadrant, we subdivide that quadrant into subquadrants, and so on. This particular splitting technique has a convenient representation in the form of a so-called quadtree (that is, a tree in which nodes have exactly four descendants), as illustrated in Fig. 7.3. Note that the root of the tree corresponds to the entire image and that each node corresponds to a subdivision. In this case, only R_4 was subdivided further.
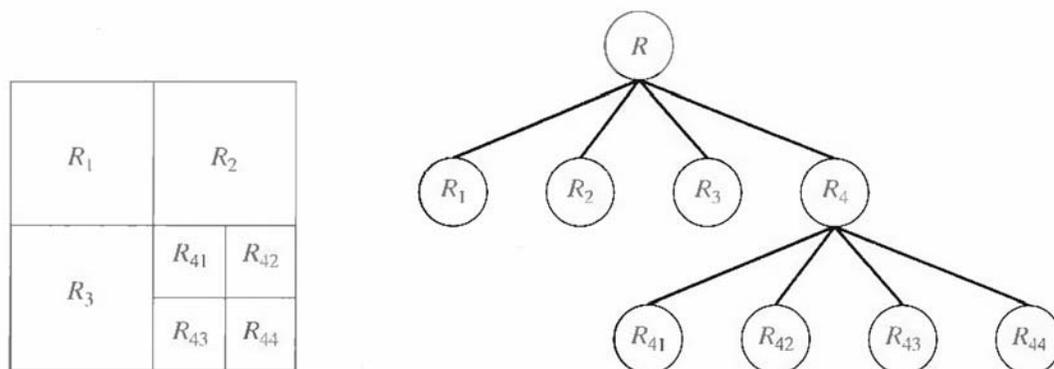
Fig. 7.3 (a) Partitioned image (b) Corresponding quadtree.

If only splitting were used, the final partition likely would contain adjacent regions with identical properties. This drawback may be remedied by allowing merging, as well as splitting. Satisfying the constraints, requires merging only adjacent regions whose combined pixels satisfy the predicate P. That is, two adjacent regions R_j and R_k are merged only if $P(R_j \cup R_k)$ = TRUE.

The preceding discussion may be summarized by the following procedure, in which, at any step we

1. Split into four disjoint quadrants any region R_i, for which $P(R_i)$ = FALSE.

2. Merge any adjacent regions R_j and R_k for which $P(R_j \cup R_k)$ = TRUE.

3. Stop when no further merging or splitting is possible.

Several variations of the preceding basic theme are possible. For example, one possibility is to split the image initially into a set of blocks. Further splitting is carried out as described previously, but merging is initially limited to groups of four blocks that are descendants in the quadtree representation and that satisfy the predicate P. When no further mergings of this type are possible, the procedure is terminated by one final merging of regions satisfying step 2. At this point, the merged regions may be of different sizes. The principal advantage of this approach is that it uses the same quadtree for splitting and merging, until the final merging step.

Image Representation

- Gaussian pyramids
- Laplacian Pyramids
- Wavelet Pyramids
- Applications

Image Pyramids

Image features at different resolutions require filters at different scales.

Edges (derivatives):

$f(x)$

$f^{'}(x)$
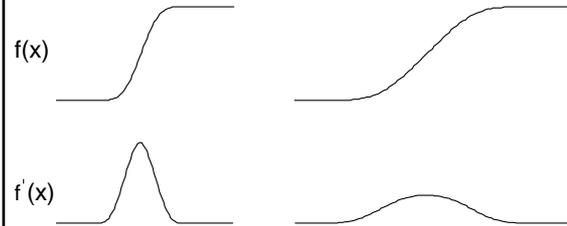
Image Pyramids

Image Pyramid = Hierarchical representation of an image

Low Resolution   No details in image - (blurred image) low frequencies

High Resolution   Details in image - low+high frequencies
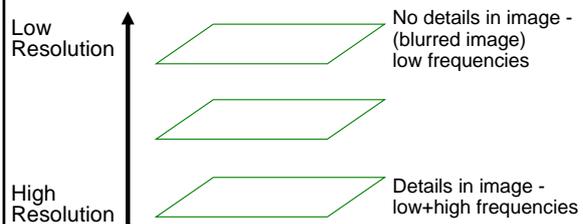
A collection of images at different resolutions.

Image pyramids

- Gaussian Pyramids
- Laplacian Pyramids
- Wavelet/QMF
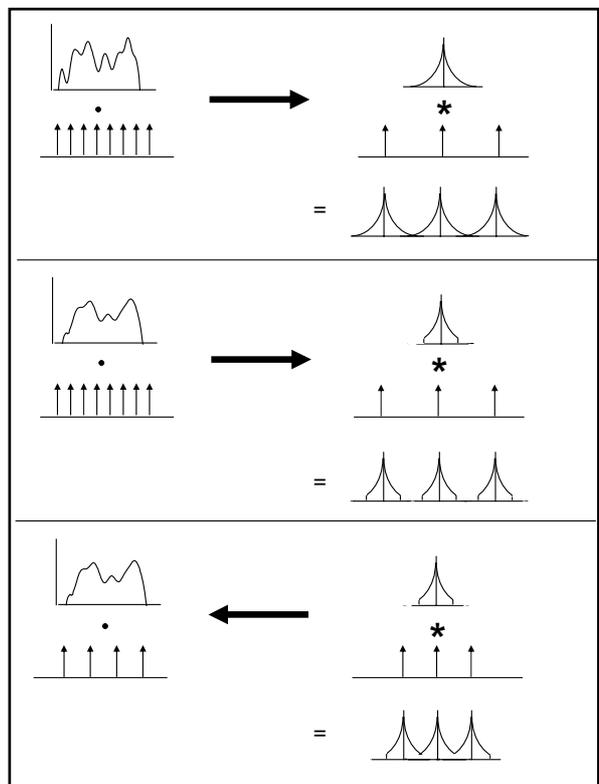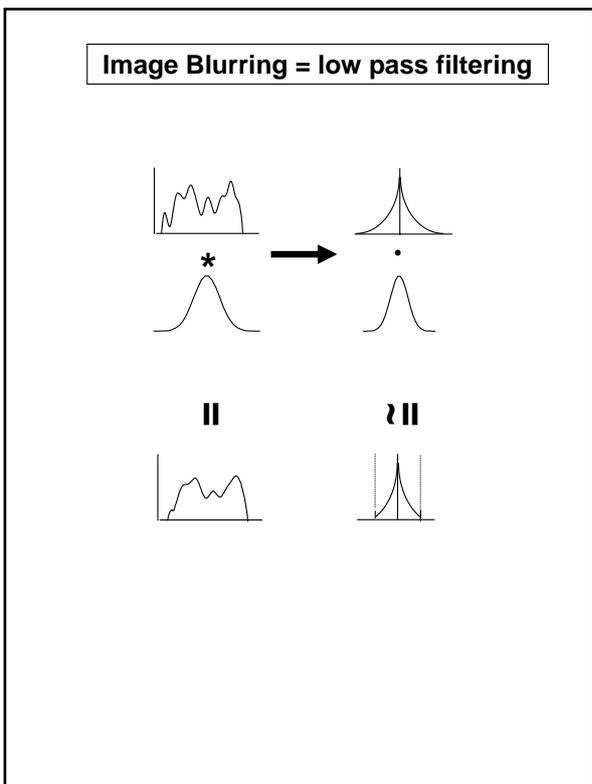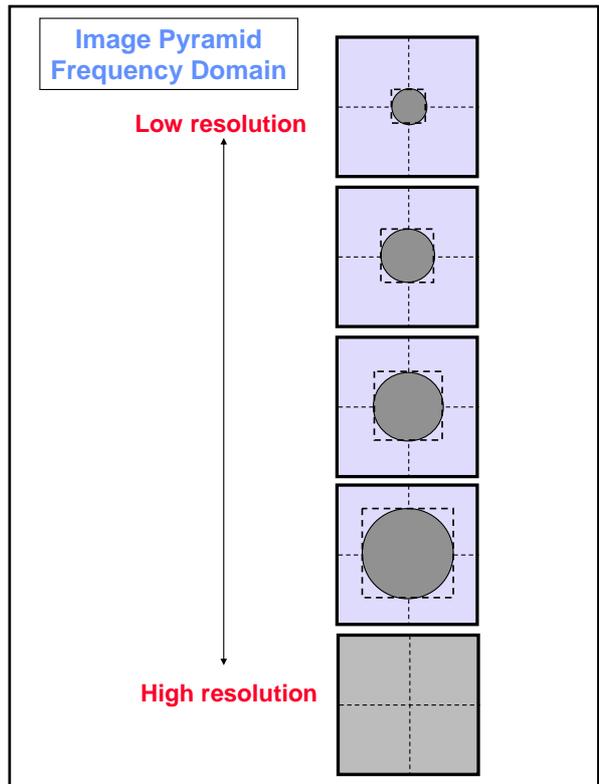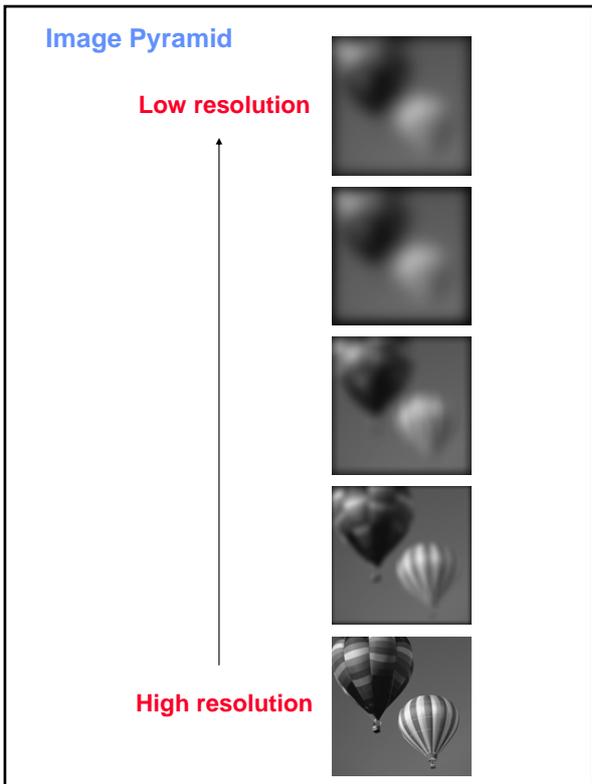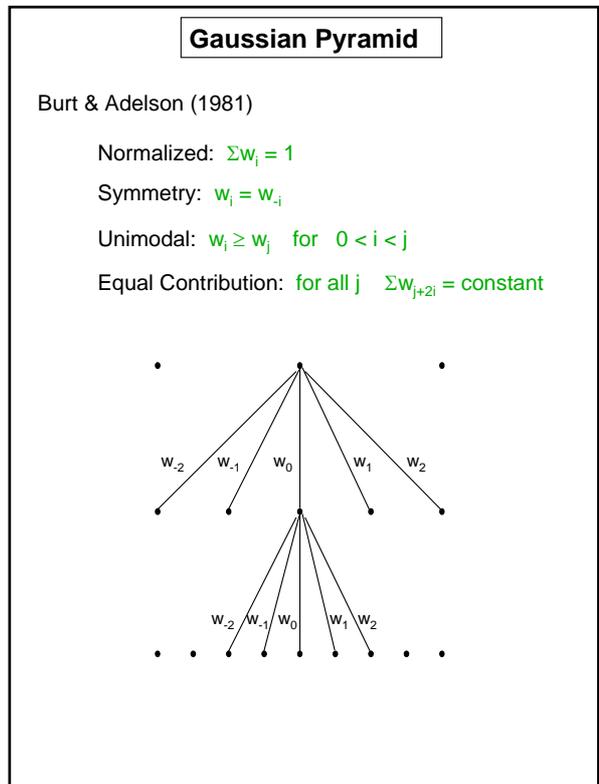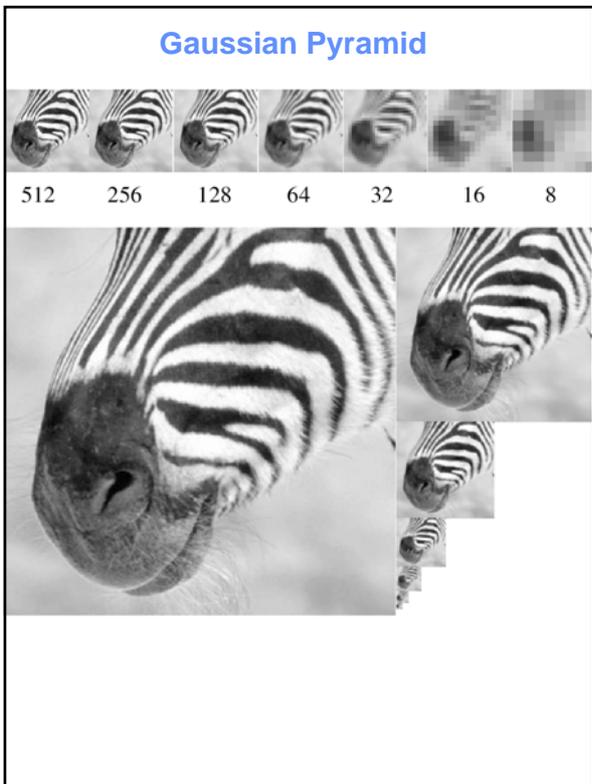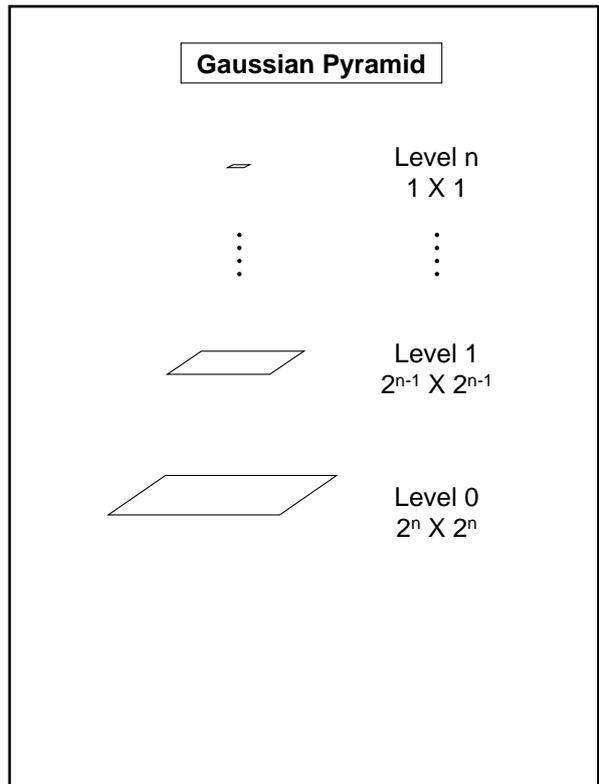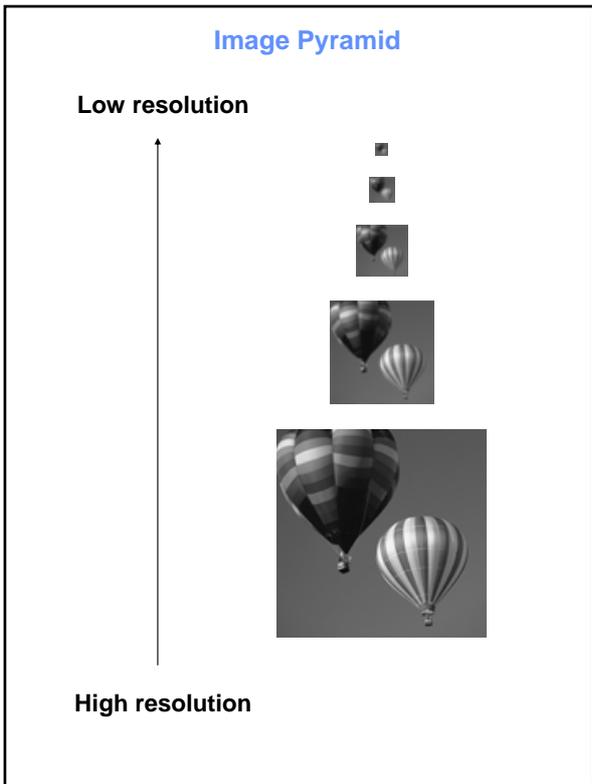
Image Pyramid

Low resolution

High resolution

**Image Pyramid
Frequency Domain**

Low resolution

High resolution

Image Blurring = low pass filtering

Image Pyramid

Low resolution

High resolution

Gaussian Pyramid

Level n
1 X 1

Level 1
$2^{n-1} \times 2^{n-1}$

Level 0
$2^n \times 2^n$

Gaussian Pyramid

512   256   128   64   32   16   8

Gaussian Pyramid

Burt & Adelson (1981)

Normalized: $\Sigma w_i = 1$

Symmetry: $w_i = w_{-i}$

Unimodal: $w_i \geq w_j$ for $0 < i < j$

Equal Contribution: for all j Σw_{j+2i} = constant

w_{-2} w_{-1} w_0 w_1 w_2

w_{-2} w_{-1} w_0 w_1 w_2

Gaussian Pyramid

Burt & Adelson (1981)

Normalized: $\Sigma w_i = 1$

Symmetry: $w_i = w_{-i}$

Unimodal: $w_i \geq w_j$ for $0 < i < j$
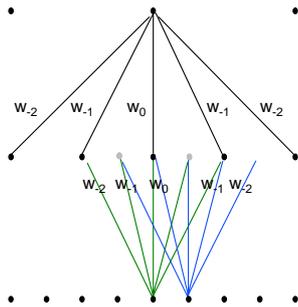
Equal Contribution: for all j $\Sigma w_{j+2i} = $ constant

Gaussian Pyramid

Burt & Adelson (1981)

Normalized: $\Sigma w_i = 1$

Symmetry: $w_i = w_{-i}$

Unimodal: $w_i \geq w_j$ for $0 < i < j$

Equal Contribution: for all j $\Sigma w_{j+2i} = $ constant
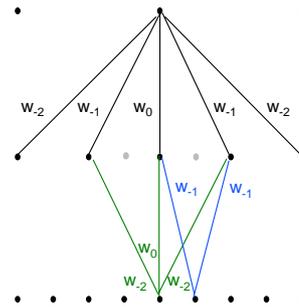
Gaussian Pyramid

Burt & Adelson (1981)
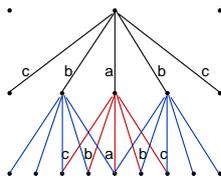
$a + 2b + 2c = 1$
$a + 2c = 2b$

$a > 0.25$
$b = 0.25$
$c = 0.25 - a/2$

For a = 0.4 most similar to a Gaussian filter

g = [0.05 0.25 0.4 0.25 0.05]

low_pass_filter = g' * g =

$$
\begin{bmatrix}
0.0025 & 0.0125 & 0.0200 & 0.0125 & 0.0025 \\
0.0125 & 0.0625 & 0.1000 & 0.0625 & 0.0125 \\
0.0200 & 0.1000 & 0.1600 & 0.1000 & 0.0200 \\
0.0125 & 0.0625 & 0.1000 & 0.0625 & 0.0125 \\
0.0025 & 0.0125 & 0.0200 & 0.0125 & 0.0025
\end{bmatrix}
$$
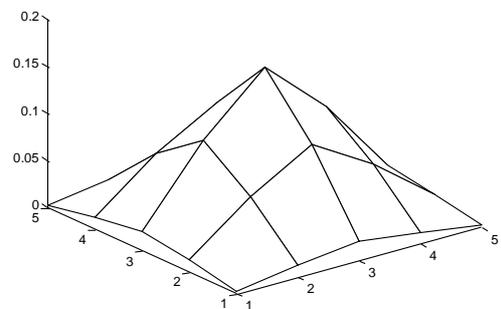
Gaussian Pyramid - Computational Aspects

Memory:

$$2^N \times 2^N (1 + 1/4 + 1/16 + \dots) = 2^N \times 2^N * 4/3$$

Computation:

Level i can be computed with a single convolution with filter: $h_i = \underbrace{g * g * g * \dots}_{i \text{ times}}$

Example:

$h_2 = $ g $*$ g $=$

MultiScale Pattern Matching

Option 1:
Scale target and search for each in image.
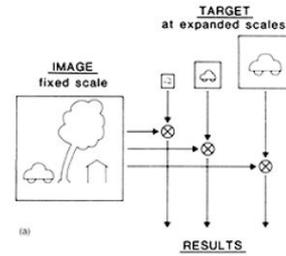
Option 2:
Search for original target in image pyramid.
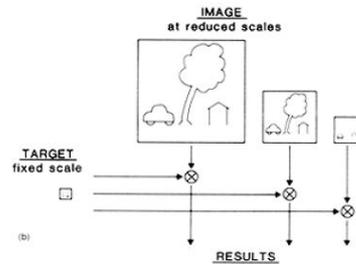
Hierarchical Pattern Matching
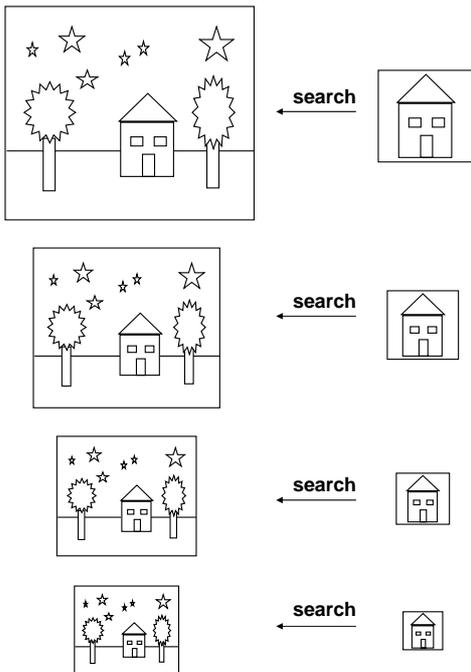
Pattern matching using Pyramids - Example

image	pattern	correlation
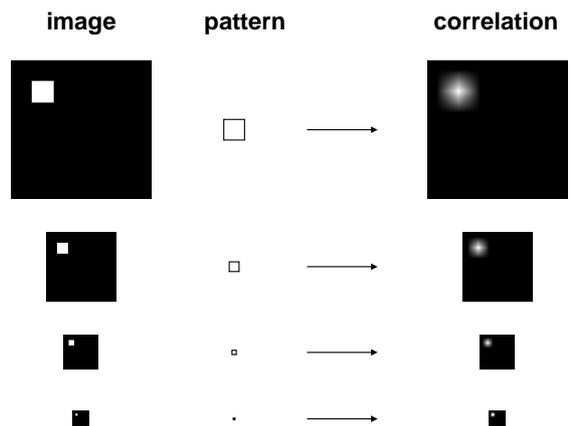
Image pyramids

- Gaussian Pyramids
- **Laplacian Pyramids**
- Wavelet/QMF

Laplacian Pyramid

Motivation = Compression, redundancy removal. compression rates are higher for predictable values. e.g. values around 0.

$G_0, G_1, ...$ = the levels of a Gaussian Pyramid.

Predict level G_l from level G_{l+1} by Expanding G_{l+1} to G'_l
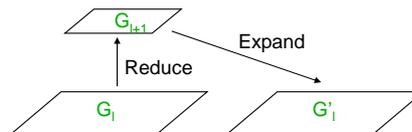
Denote by L_l the error in prediction:

$$L_l = G_l - G'_l$$

$L_0, L_1, ...$ = the levels of a Laplacian Pyramid.

What does blurring take away?

original

What does blurring take away?
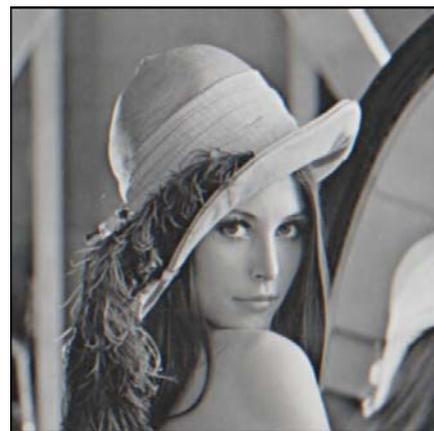
smoothed (5x5 Gaussian)

What does blurring take away?

smoothed – original

Laplacian Pyramid

Gaussian Pyramid

Laplacian Pyramid

Frequency Domain

Gaussian Pyramid
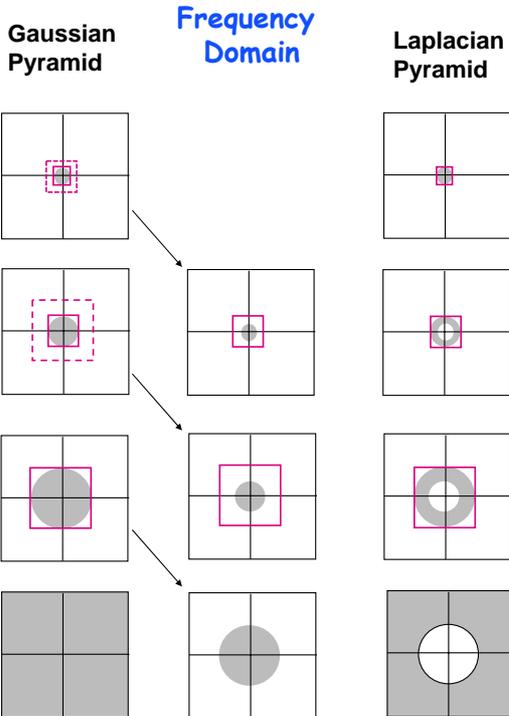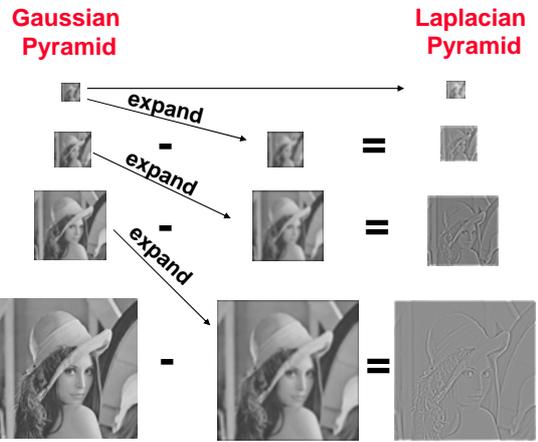
Laplacian Pyramid

Laplace Pyramid - No scaling

from: B.Freeman

| 512 | 256 | 128 | 64 | 32 | 16 | 8 |

Reconstruction of the original image from the Laplacian Pyramid

Laplacian Pyramid

$$G_l = L_l + G'_l$$

expand

expand

expand

= **Original Image**
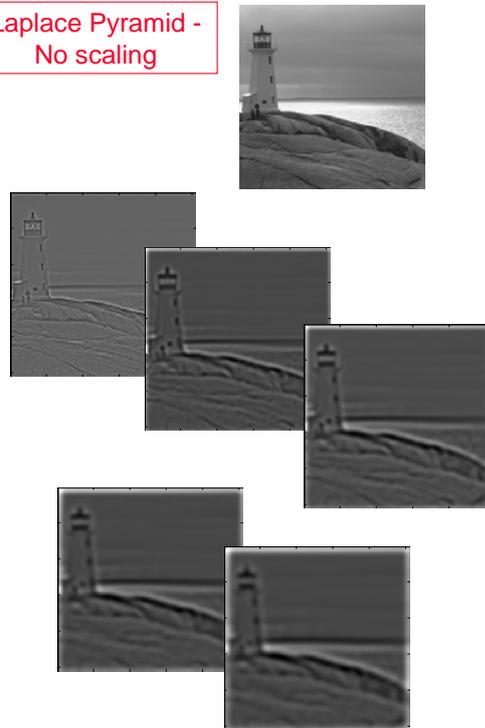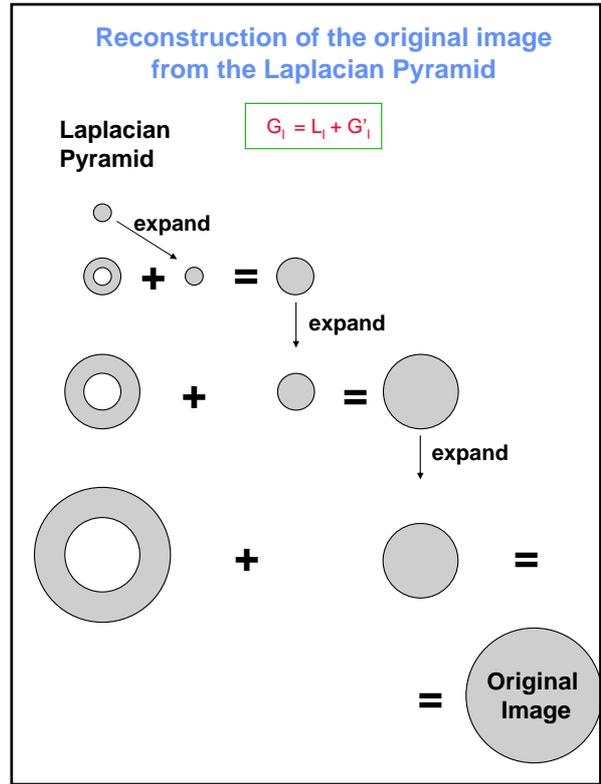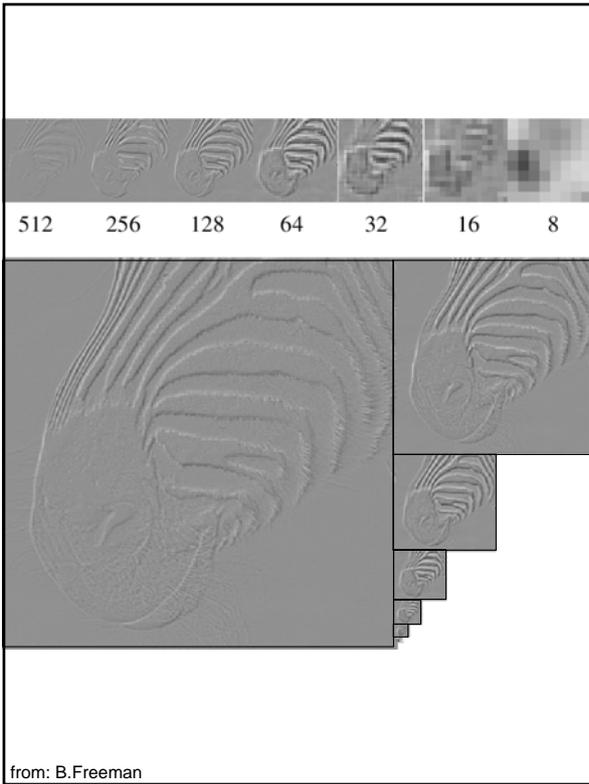
Laplacian Pyramid - Computational Aspects

Memory:

$$2^N \times 2^N (1 + 1/4 + 1/16 + \ldots) = 2^N \times 2^N * 4/3$$

However coefficients are highly compressible.

Computation:

L_i can be computed from G_0 with a single convolution with filter: $k_i = h_{i-1} - h_i$

h_{i-1} $-$ h_i $=$ k_i

k_1 k_2 k_3

Image Mosaicing

Registration
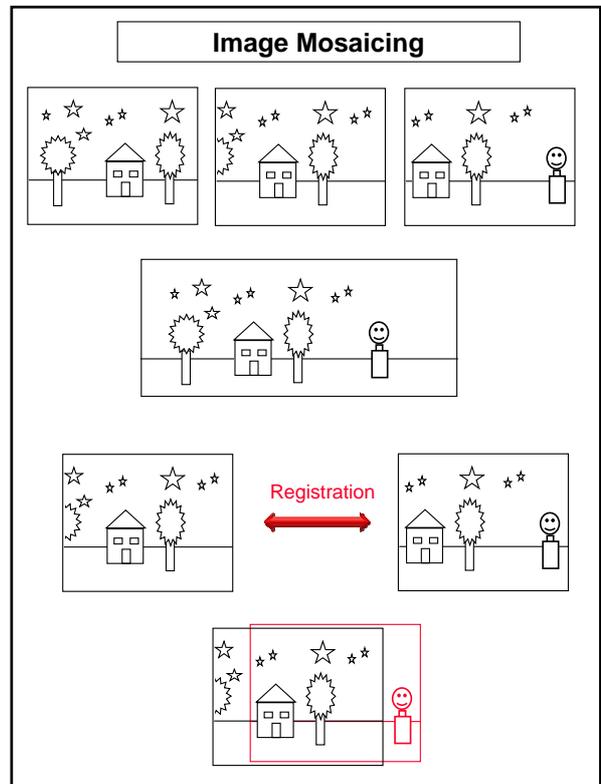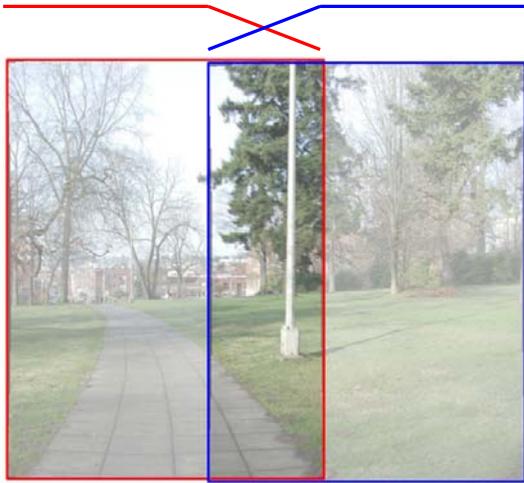
Image Blending

Blending

Multiresolution Spline

When splining two images, transition from one image to the other should behave:

High Frequencies

Middle Frequencies

Low Frequencies

Multiresolution Spline

High Frequencies

Middle Frequencies

Low Frequencies

Multiresolution Spline - Example

Left Image

Right Image

Left + Right

Narrow Transition

Wide Transition
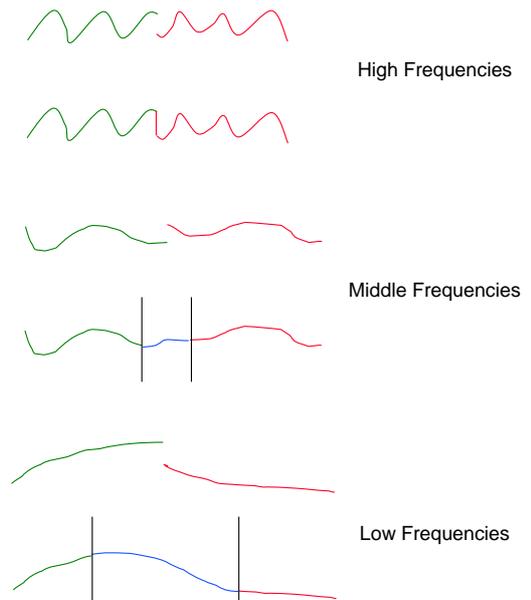
(Burt & Adelson)

Multiresolustion Spline - Using Laplacian Pyramid
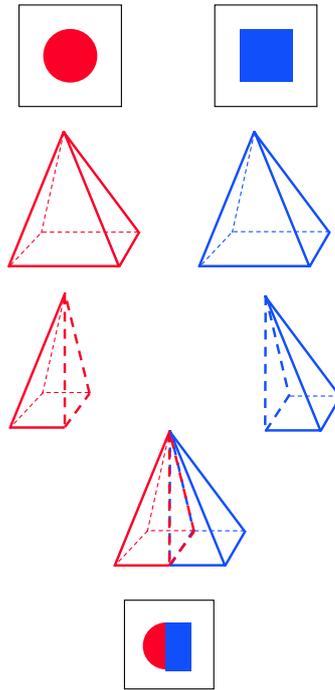
Multiresolution Spline - Example

Left Image

Right Image

Left + Right

Narrow Transition

Wide Transition

Multiresolution Spline

(Burt & Adelson)

Multiresolution Spline - Example

Original - Left

Original - Right

Glued
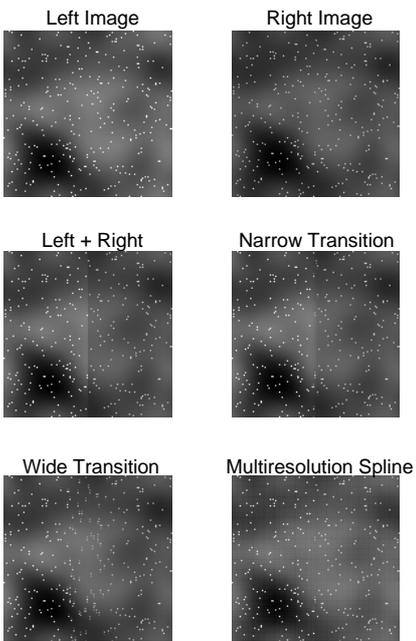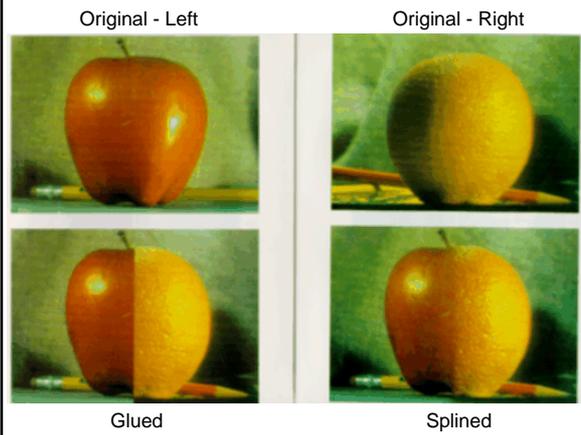
Splined
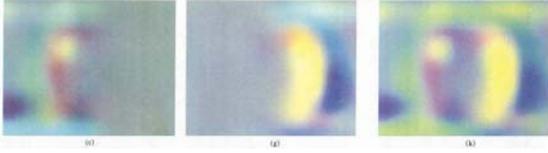
laplacian level 4

laplacian level 2
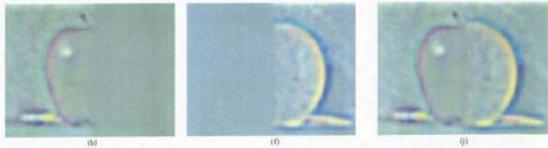
laplacian level 0
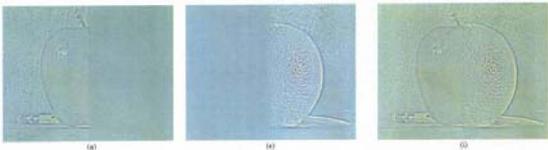
left pyramid   right pyramid   blended pyramid

Multiresolution Spline
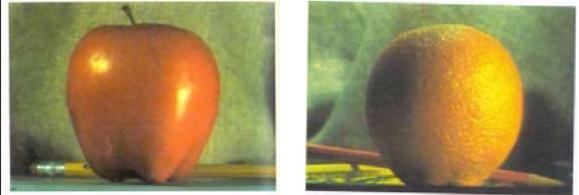
Multiresolution Spline - Example

Multi-Res. Blending
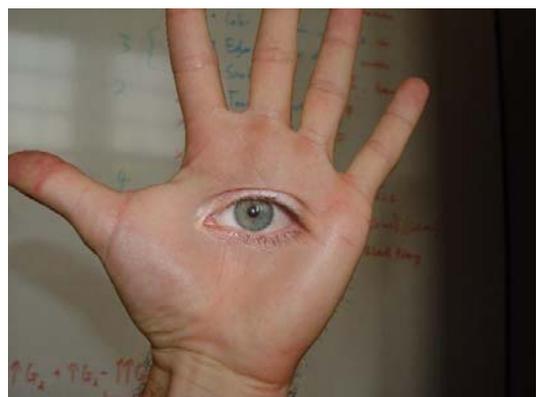
© prof. dmartin

Image pyramids

- Gaussian Pyramids
- Laplacian Pyramids
- **Wavelet/QMF**

What is a good representation for image analysis?

- Pixel domain representation tells you "where" (pixel location), but not "what".
  - In space, this representation is too localized
- Fourier transform domain tells you "what" (textural properties), but not "where".
  - In space, this representation is too spread out.
- Want an image representation that gives you a local description of image events—what is happening where.
  - That representation might be "just right".

Space-Frequency Tiling
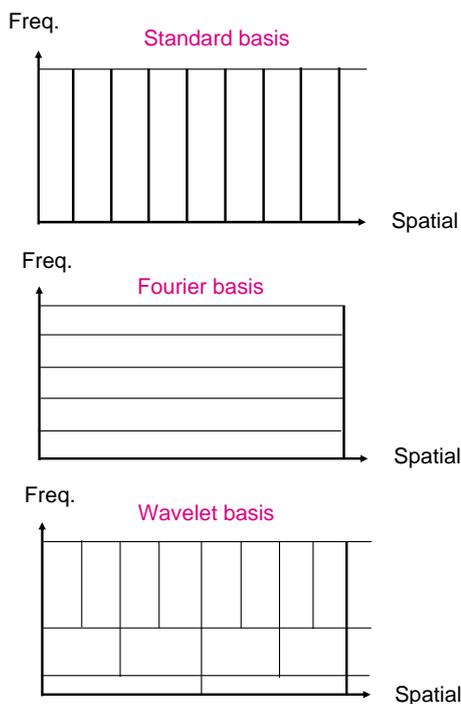
Space-Frequency Tiling
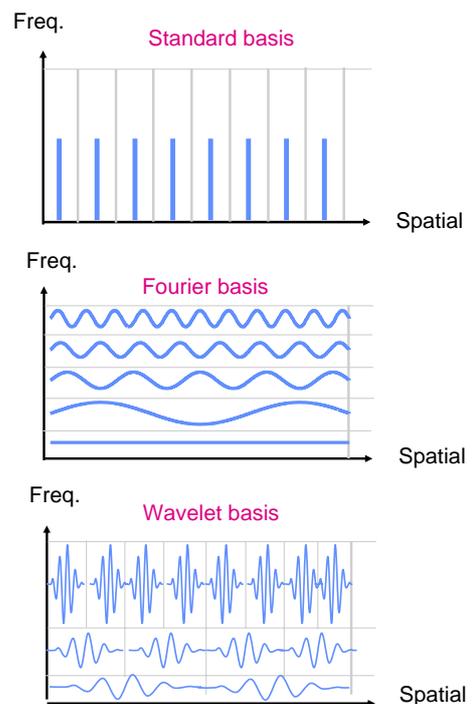
Various Wavelet basis
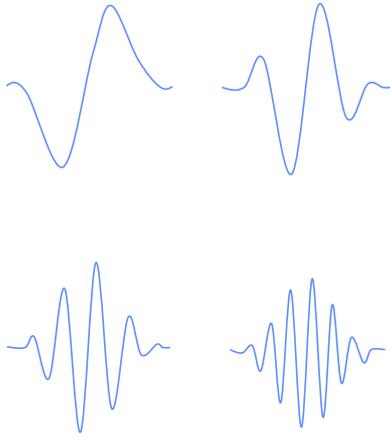
Wavelet - Frequency domain

Wavelet bands are split recursively

image

H   L

H   L

H   L
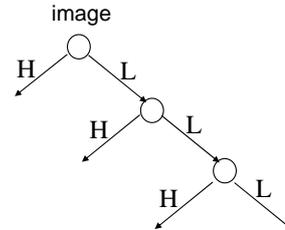
Wavelet - Frequency domain

Wavelet decomposition - 2D

Frequency domain

Horizontal high pass

Vertical high pass

Horizontal low pass

Vertical low pass
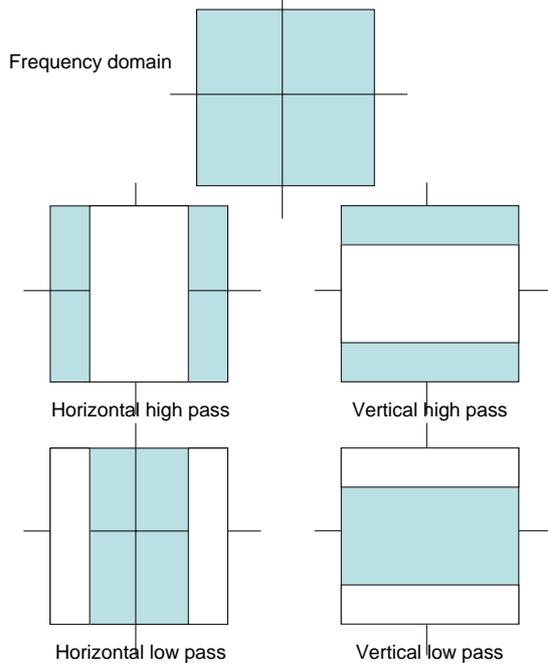
Wavelet - Frequency domain

Apply the wavelet transform separably in both dimensions

Horizontal high pass,
vertical high pass

Horizontal high pass,
vertical low-pass

Horizontal low pass,
vertical high-pass

Horizontal low pass,
Vertical low-pass
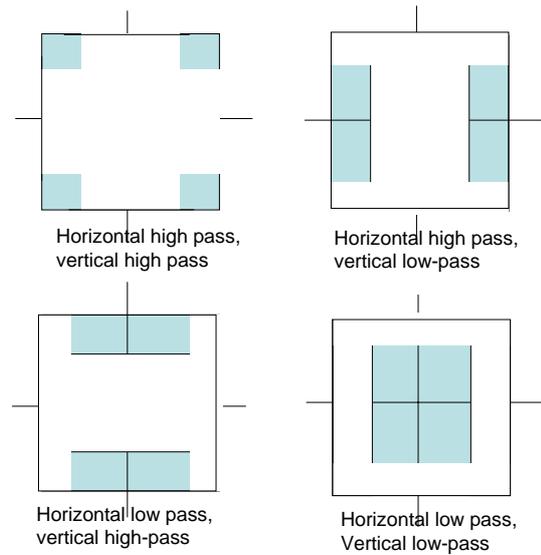
- Splitting can be applied recursively:
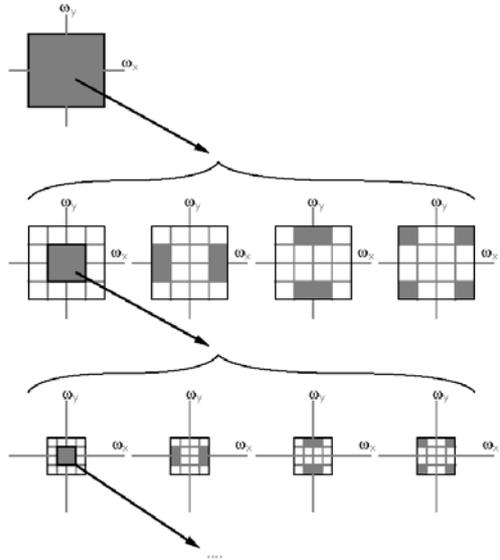
Pyramids in Frequency Domain

Gaussian Pyramid   **Laplacian Pyramid**

Wavelet Decomposition

Fourier Space
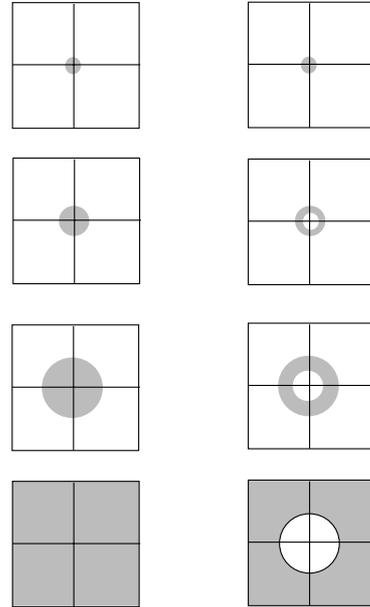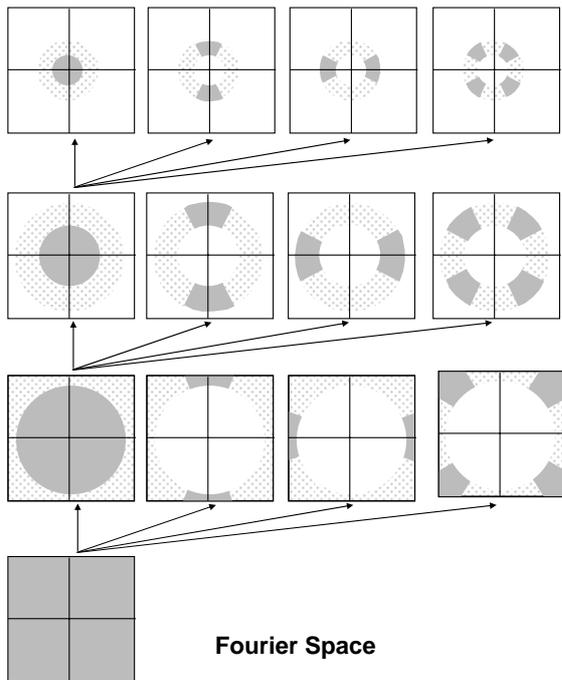
Wavelet Transform - Step By Step Example
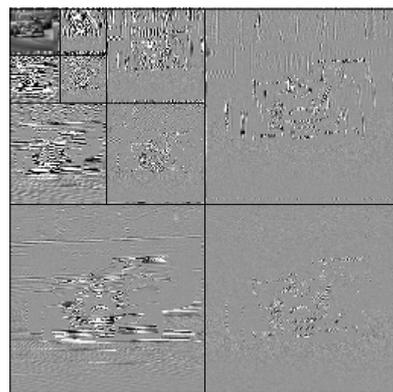
Wavelet Transform - Example
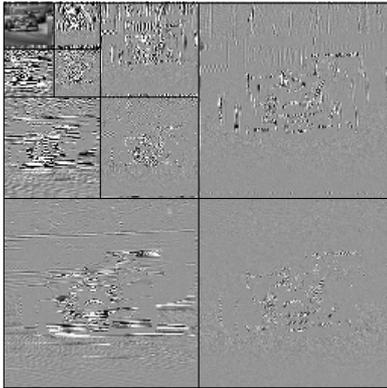
Wavelet Transform - Example
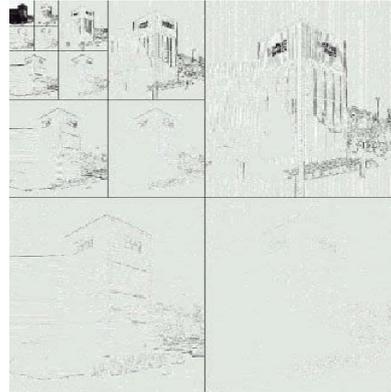
Application: Wavelet Shrinkage Denoising

Noisy image
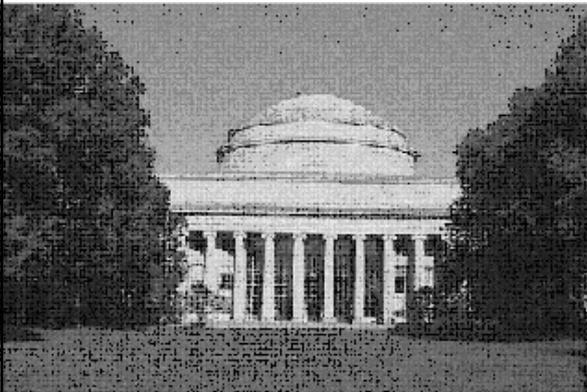
Clean image

Range [0, 255]
Dims [394, 599]

Wavelet Shrinkage Denoising

Noisy image

Wavelet coefficient Histogram

Clean image

Wavelet coefficient Histogram
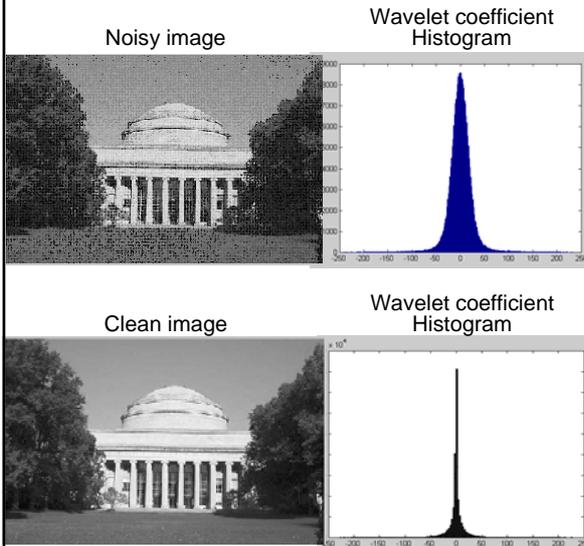
Get top histogram but want to get bottom histogram.

From: B. Freeman

Wavelet Shrinkage Denoising

For every Wavelet Band define Shrinkage function:

New Coefficient Value

Wavelet Coefficient Value
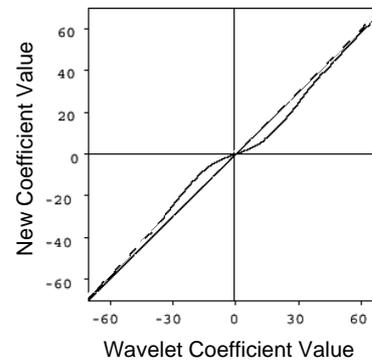
From: B. Freeman

Wavelet Shrinkage Pipe-line

I → Transform W → Mapping functions x^i_w / y^i_w → Inverse Transform W^T → I_{clean}
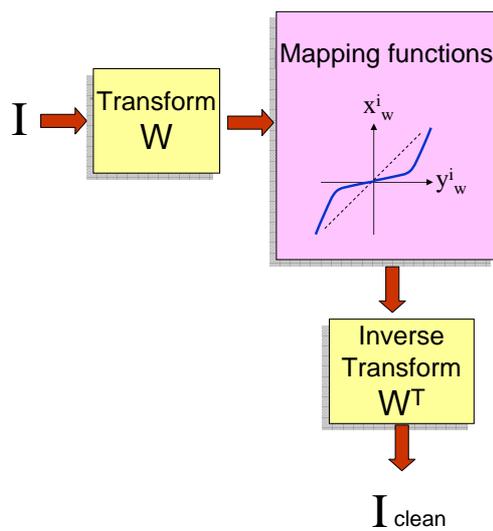
More results

More results

Image Pyramids - Comparison

Image pyramid levels = Filter then sample.

Filters:

Gaussian Pyramid

Laplacian Pyramid
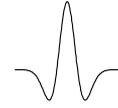
Wavelet Pyramid
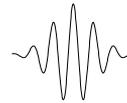
Image Linear Transforms

Transform	Basis	Characteristics
Delta	Standard	Localized in space Not localized in Frequency
Fourier	Sines+Cosines	Not localized in space Localized in Frequency
Wavelet Pyramid	Wavelet Filters	Localized in space Localized in Frequency

Delta	Fourier	Wavelet
space	space	space
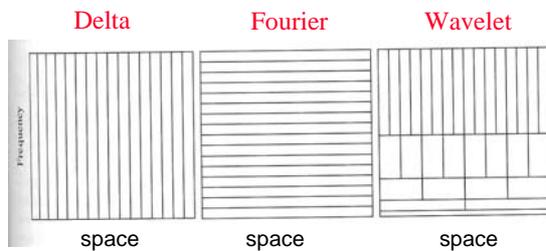
Convolution and Transforms in matrix notation (1D case)

$$\vec{F} = U\vec{f}$$

transformed image

Basis vectors (Fourier, Wavelet, etc)

Vectorized image

Fourier Transform

$$\Bigg| \Bigg. = \boxed{} * \Bigg| \Bigg.$$

Fourier transform   Fourier bases   pixel domain image
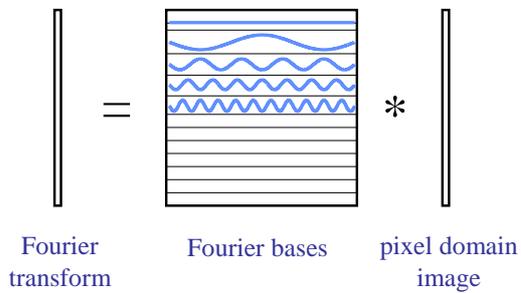
Fourier bases are global: each transform coefficient depends on all pixel locations.

Transform in matrix notation (1D case)

Forward Transform:

$$\vec{F} = U\vec{f}$$

transformed image

Basis vectors (Fourier, Wavelet, etc)

Vectorized image

Inverse Transform:

$$U^{-1}\vec{F} = \vec{f}$$

Basis vectors

transformed image

Vectorized image

Inverse Fourier Transform

$$\boxed{} * \Bigg| \Bigg. = \Bigg| \Bigg.$$

Fourier bases   Fourier transform   pixel domain image
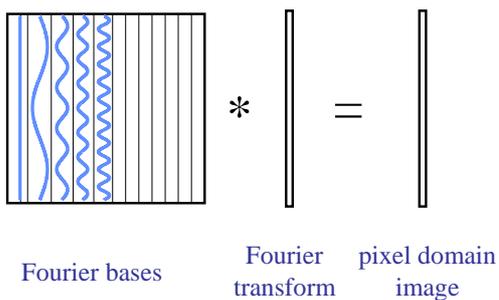
Every image pixel is a linear combination of the Fourier basis weighted by the coefficient.

Note that if U is orthonormal basis then $U^{-1} = U^t$

Convolution

$$\Bigg| \Bigg. = \boxed{} * \Bigg| \Bigg.$$

Convolution Result   Circular Matrix of Filter Kernels   pixel image

Pyramid =
Convolution + Sampling

Convolution   Matrix of   pixel image
Result   Filter Kernels
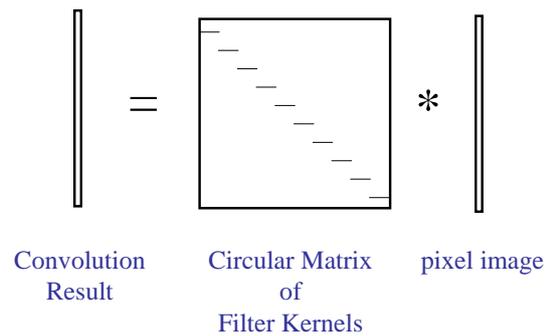
From: B. Freeman

Pyramid =
Convolution + Sampling

Pyramid Level 1
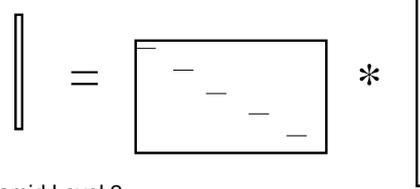
Pyramid Level 2
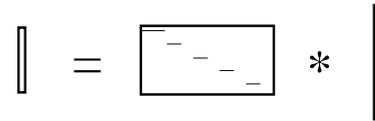
Pyramid Level 2
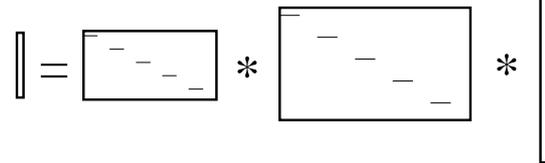
Pyramid =
Convolution + Sampling

Pyramid Level 2
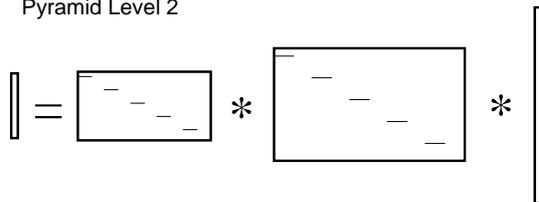
Pyramid Level 2
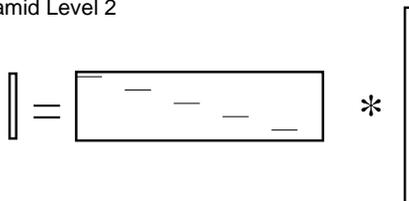
Pyramid as Matrix
Computation - Example

U1 =

1	4	6	4	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	1	4	6	4	1	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	1	4	6	4	1	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	1	4	6	4	1	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	1	4	6	4	1	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	1	4	6	4	1	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	1	4	6	4	1	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	4	6	4	1	0

- Next pyramid level

U2 =

1	4	6	4	1	0	0	0
0	0	1	4	6	4	1	0
0	0	0	0	1	4	6	4
0	0	0	0	0	0	1	4

- The combined effect of the two pyramid levels

U2 * U1 =

1	4	10	20	31	40	44	40	31	20	10	4	1	0	0	0	0	0	0	0
0	0	0	0	1	4	10	20	31	40	44	40	31	20	10	4	1	0	0	0
0	0	0	0	0	0	0	0	1	4	10	20	31	40	44	40	30	16	4	0
0	0	0	0	0	0	0	0	0	0	0	0	0	1	4	10	20	25	16	4	0

from: B.Freeman

Gaussian Pyramid
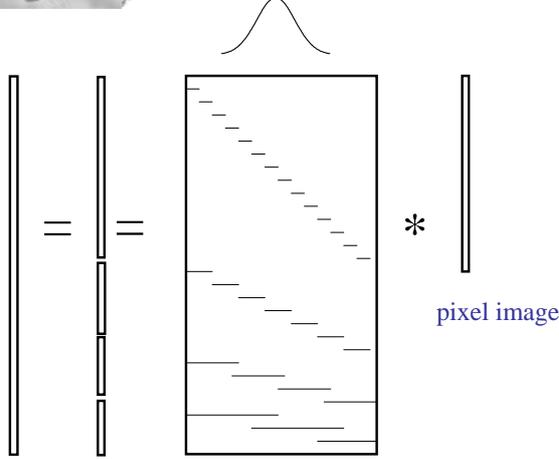
= ‖ = [filter matrix] * [pixel image]

pixel image

Gaussian pyramid

Overcomplete representation. Low-pass filters, sampled appropriately for their blur.

From: B. Freeman

Laplacian Pyramid
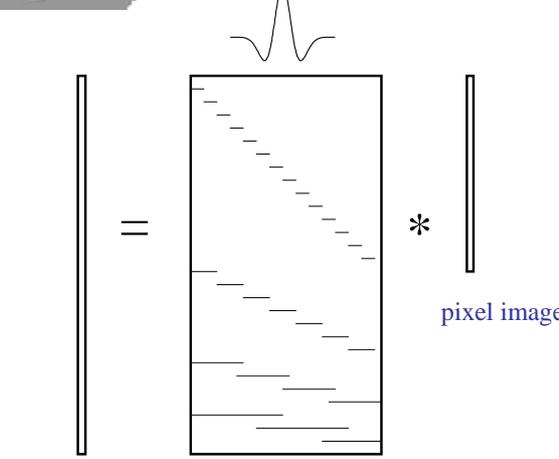
= [filter matrix] * [pixel image]

pixel image

Laplacian pyramid

Overcomplete representation. Transformed pixels represent bandpassed image information.

From: B. Freeman

Wavelet Transform

= [filter matrix] * [pixel image]
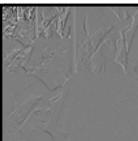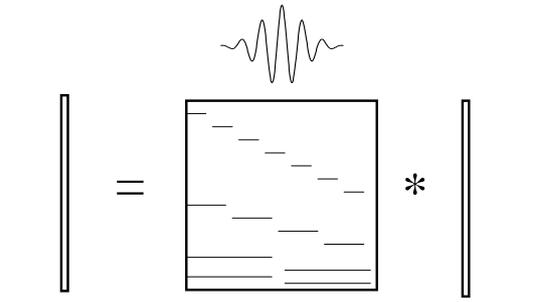
Wavelet pyramid

pixel image

Ortho-normal transform (like Fourier transform), but with localized basis functions.

From: B. Freeman

The End

1.Define image compression. Explain about the redundancies in a digital image.

The term data compression refers to the process of reducing the amount of data required to represent a given quantity of information. A clear distinction must be made between data and information. They are not synonymous. In fact, data are the means by which information is conveyed. Various amounts of data may be used to represent the same amount of information. Such might be the case, for example, if a long-winded individual and someone who is short and to the point were to relate the same story. Here, the information of interest is the story; words are the data used to relate the information. If the two individuals use a different number of words to tell the same basic story, two different versions of the story are created, and at least one includes nonessential data. That is, it contains data (or words) that either provide no relevant information or simply restate that which is already known. It is thus said to contain data redundancy.

Data redundancy is a central issue in digital image compression. It is not an abstract concept but a mathematically quantifiable entity. If n_1 and n_2 denote the number of information-carrying units in two data sets that represent the same information, the relative data redundancy R_D of the first data set (the one characterized by n_1) can be defined as

$$R_D = 1 - \frac{1}{C_R}$$

where C_R , commonly called the compression ratio, is

$$C_R = \frac{n_1}{n_2}.$$

For the case $n_2 = n_1, C_R = 1$ and $R_D = 0$, indicating that (relative to the second data set) the first representation of the information contains no redundant data. When $n_2 \ll n_1$, $C_R \rightarrow \infty$ and $R_D \rightarrow 1$, implying significant compression and highly redundant data. Finally, when $n_2 \gg n_1$, $C_R \rightarrow 0$ and $R_D \rightarrow \infty$, indicating that the second data set contains much more data than the original representation. This, of course, is the normally undesirable case of data expansion. In general, C_R and R_D lie in the open intervals $(0,\infty)$ and $(-\infty, 1)$, respectively. A practical compression ratio, such as 10 (or 10:1), means that the first data set has 10 information carrying units (say, bits) for every 1 unit in the second or compressed data set. The corresponding redundancy of 0.9 implies that 90% of the data in the first data set is redundant.

In digital image compression, three basic data redundancies can be identified and exploited: **coding redundancy**, **interpixel redundancy**, and **psychovisual redundancy**. Data compression is achieved when one or more of these redundancies are reduced or eliminated.

Coding Redundancy:

In this, we utilize formulation to show how the gray-level histogram of an image also can provide a great deal of insight into the construction of codes to reduce the amount of data used to represent it.

Let us assume, once again, that a discrete random variable r_k in the interval [0, 1] represents the gray levels of an image and that each r_k occurs with probability $p_r(r_k)$.

$$p_r(r_k) = \frac{n_k}{n} \qquad k = 0, 1, 2, \dots, L - 1$$

where L is the number of gray levels, n_k is the number of times that the kth gray level appears in the image, and n is the total number of pixels in the image. If the number of bits used to represent each value of r_k is $l(r_k)$, then the average number of bits required to represent each pixel is

$$L_{avg} = \sum_{k=0}^{L-1} l(r_k) p_r(r_k).$$

That is, the average length of the code words assigned to the various gray-level values is found by summing the product of the number of bits used to represent each gray level and the probability that the gray level occurs. Thus the total number of bits required to code an M X N image is MNL_{avg}.

Interpixel Redundancy:

Consider the images shown in Figs. 1.1(a) and (b). As Figs. 1.1(c) and (d) show, these images have virtually identical histograms. Note also that both histograms are trimodal, indicating the presence of three dominant ranges of gray-level values. Because the gray levels in these images are not equally probable, variable-length coding can be used to reduce the coding redundancy that would result from a straight or natural binary encoding of their pixels. The coding process, however, would not alter the level of correlation between the pixels within the images. In other words, the codes used to represent the gray levels of each image have nothing to do with the correlation between pixels. These correlations result from the structural or geometric relationships between the objects in the image.
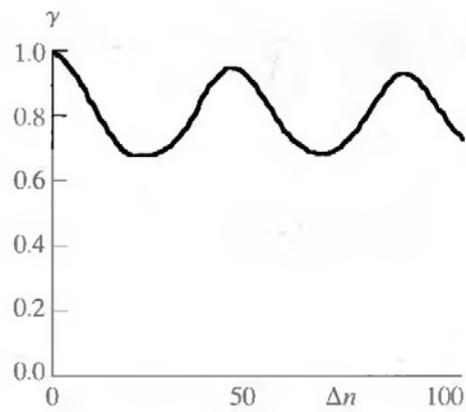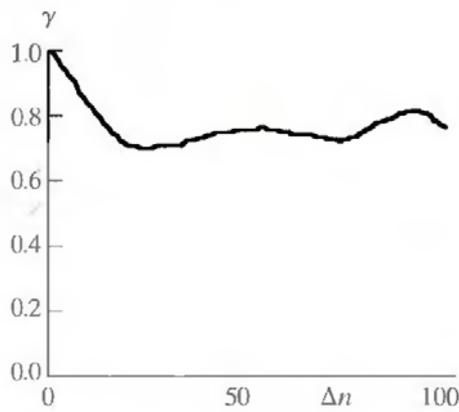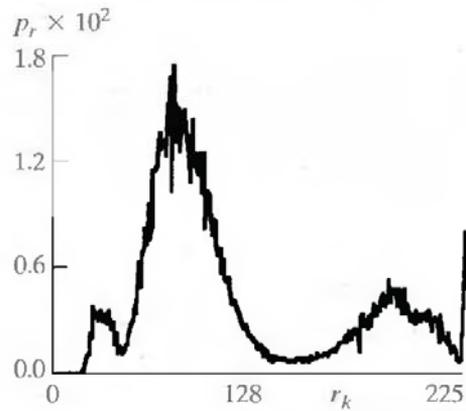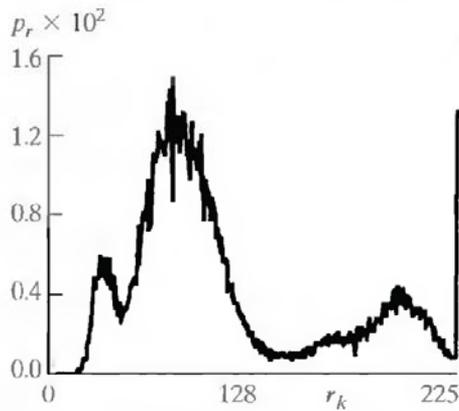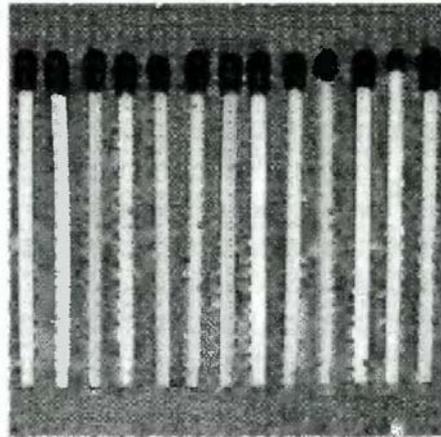
Fig.1.1 Two images and their gray-level histograms and normalized autocorrelation coefficients along one line.

Figures 1.1(e) and (f) show the respective autocorrelation coefficients computed along one line of each image.

$$\gamma(\Delta n) = \frac{A(\Delta n)}{A(0)}$$

where

$$A(\Delta n) = \frac{1}{N - \Delta n} \sum_{y=0}^{N-1-\Delta n} f(x, y) f(x, y + \Delta n).$$

The scaling factor in Eq. above accounts for the varying number of sum terms that arise for each integer value of Δn. Of course, Δn must be strictly less than N, the number of pixels on a line. The variable x is the coordinate of the line used in the computation. Note the dramatic difference between the shape of the functions shown in Figs. 1.1(e) and (f). Their shapes can be qualitatively related to the structure in the images in Figs. 1.1(a) and (b).This relationship is particularly noticeable in Fig. 1.1 (f), where the high correlation between pixels separated by 45 and 90 samples can be directly related to the spacing between the vertically oriented matches of Fig. 1.1(b). In addition, the adjacent pixels of both images are highly correlated. When Δn is 1, γ is 0.9922 and 0.9928 for the images of Figs. 1.1 (a) and (b), respectively. These values are typical of most properly sampled television images.

These illustrations reflect another important form of data redundancy—one directly related to the interpixel correlations within an image. Because the value of any given pixel can be reasonably predicted from the value of its neighbors, the information carried by individual pixels is relatively small. Much of the visual contribution of a single pixel to an image is redundant; it could have been guessed on the basis of the values of its neighbors. A variety of names, including spatial redundancy, geometric redundancy, and interframe redundancy, have been coined to refer to these interpixel dependencies. We use the term interpixel redundancy to encompass them all.

In order to reduce the interpixel redundancies in an image, the 2-D pixel array normally used for human viewing and interpretation must be transformed into a more efficient (but usually "nonvisual") format. For example, the differences between adjacent pixels can be used to represent an image. Transformations of this type (that is, those that remove interpixel redundancy) are referred to as mappings. They are called reversible mappings if the original image elements can be reconstructed from the transformed data set.

 Psychovisual Redundancy:

The brightness of a region, as perceived by the eye, depends on factors other than simply the light reflected by the region. For example, intensity variations (Mach bands) can be perceived in an area of constant intensity. Such phenomena result from the fact that the eye does not respond with equal sensitivity to all visual information. Certain information simply has less relative importance than other information in normal visual processing. This information is said to be psychovisually redundant. It can be eliminated without significantly impairing the quality of image perception.

That psychovisual redundancies exist should not come as a surprise, because human perception of the information in an image normally does not involve quantitative analysis of every pixel value in the image. In general, an observer searches for distinguishing features such as edges or textural regions and mentally combines them into recognizable groupings. The brain then correlates these groupings with prior knowledge in order to complete the image interpretation process. Psychovisual redundancy is fundamentally different from the redundancies discussed earlier. Unlike coding and interpixel redundancy, psychovisual redundancy is associated with real or quantifiable visual information. Its elimination is possible only because the information itself is not essential for normal visual processing. Since the elimination of psychovisually redundant data results in a loss of quantitative information, it is commonly referred to as quantization.

This terminology is consistent with normal usage of the word, which generally means the mapping of a broad range of input values to a limited number of output values. As it is an irreversible operation (visual information is lost), quantization results in lossy data compression.

2. Explain about fidelity criterion.

The removal of psychovisually redundant data results in a loss of real or quantitative visual information. Because information of interest may be lost, a repeatable or reproducible means of quantifying the nature and extent of information loss is highly desirable. Two general classes of criteria are used as the basis for such an assessment:

A) Objective fidelity criteria and

B) Subjective fidelity criteria.

When the level of information loss can be expressed as a function of the original or input image and the compressed and subsequently decompressed output image, it is said to be based on an objective fidelity criterion. A good example is the root-mean-square (rms) error between an input and output image. Let f(x, y) represent an input image and let f(x, y) denote an estimate or approximation of f(x, y) that results from compressing and subsequently decompressing the input. For any value of x and y, the error e(x, y) between f (x, y) and f^ (x, y) can be defined as

$$e(x, y) = \hat{f}(x, y) - f(x, y)$$

so that the total error between the two images is

$$\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}[\hat{f}(x, y) - f(x, y)]$$

where the images are of size M X N. The root-mean-square error, e_{rms}, between f(x, y) and f^(x, y) then is the square root of the squared error averaged over the M X N array, or

$$e_{rms} = \left[\frac{1}{MN}\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}[\hat{f}(x, y) - f(x, y)]^2\right]^{1/2}$$

A closely related objective fidelity criterion is the mean-square signal-to-noise ratio of the compressed-decompressed image. If f^ (x, y) is considered to be the sum of the original image f(x, y) and a noise signal e(x, y), the mean-square signal-to-noise ratio of the output image, denoted SNR_{rms}, is

$$SNR_{ms} = \frac{\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}\hat{f}(x, y)^2}{\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}[\hat{f}(x, y) - f(x, y)]^2}.$$

The rms value of the signal-to-noise ratio, denoted SNR_{rms}, is obtained by taking the square root of Eq. above.

Although objective fidelity criteria offer a simple and convenient mechanism for evaluating information loss, most decompressed images ultimately are viewed by humans. Consequently, measuring image quality by the subjective evaluations of a human observer often is more appropriate. This can be accomplished by showing a "typical" decompressed image to an appropriate cross section of viewers and averaging their evaluations. The evaluations may be made using an absolute rating scale or by means of side-by-side comparisons of f(x, y) and f^(x, y).

3. Explain about image compression models.

 Fig. 3.1 shows, a compression system consists of two distinct structural blocks: an encoder and a decoder. An input image f(x, y) is fed into the encoder, which creates a set of symbols from the input data. After transmission over the channel, the encoded representation is fed to the decoder, where a reconstructed output image f^(x, y) is generated. In general, f^(x, y) may or may not be an exact replica of f(x, y). If it is, the system is error free or information preserving; if not, some level of distortion is present in the reconstructed image. Both the encoder and decoder shown in Fig. 3.1 consist of two relatively independent functions or subblocks. The encoder is made up of a source encoder, which removes input redundancies, and a channel encoder, which increases the noise immunity of the source encoder's output. As would be expected, the decoder includes a channel decoder followed by a source decoder. If the channel between the encoder and decoder is noise free (not prone to error), the channel encoder and decoder are omitted, and the general encoder and decoder become the source encoder and decoder, respectively.
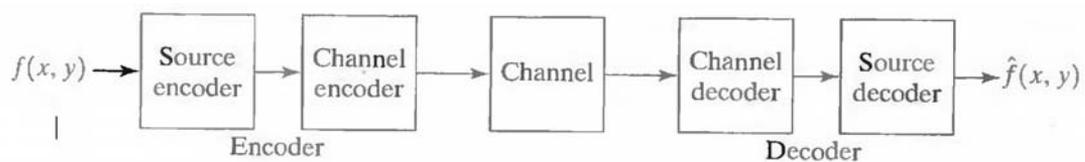
Fig.3.1 A general compression system model

 The Source Encoder and Decoder:

The source encoder is responsible for reducing or eliminating any coding, interpixel, or psychovisual redundancies in the input image. The specific application and associated fidelity requirements dictate the best encoding approach to use in any given situation. Normally, the approach can be modeled by a series of three independent operations. As Fig. 3.2 (a) shows, each operation is designed to reduce one of the three redundancies. Figure 3.2 (b) depicts the corresponding source decoder. In the first stage of the source encoding process, the mapper transforms the input data into a (usually nonvisual) format designed to reduce interpixel redundancies in the input image. This operation generally is reversible and may or may not reduce directly the amount of data required to represent the image.
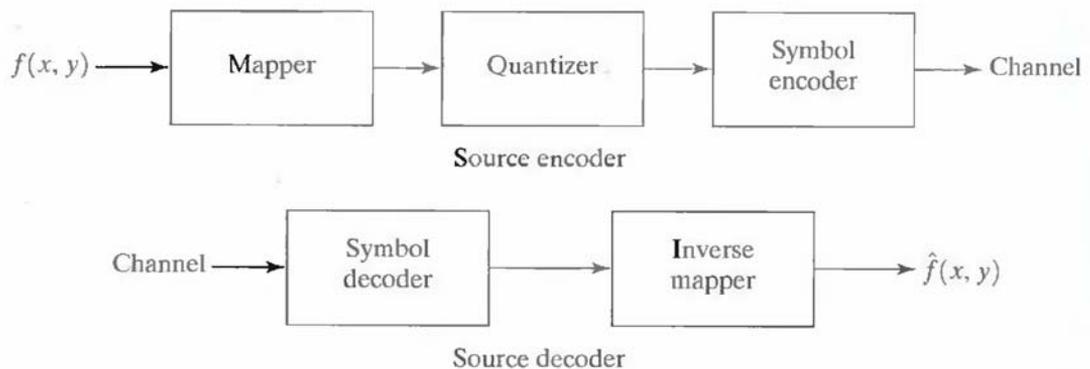
a
b

Fig.3.2 (a) Source encoder and (b) source decoder model

Run-length coding is an example of a mapping that directly results in data compression in this initial stage of the overall source encoding process. The representation of an image by a set of transform coefficients is an example of the opposite case. Here, the mapper transforms the image into an array of coefficients, making its interpixel redundancies more accessible for compression in later stages of the encoding process.

The second stage, or quantizer block in Fig. 3.2 (a), reduces the accuracy of the mapper's output in accordance with some preestablished fidelity criterion. This stage reduces the psychovisual redundancies of the input image. This operation is irreversible. Thus it must be omitted when error-free compression is desired.

In the third and final stage of the source encoding process, the symbol coder creates a fixed- or variable-length code to represent the quantizer output and maps the output in accordance with the code. The term symbol coder distinguishes this coding operation from the overall source encoding process. In most cases, a variable-length code is used to represent the mapped and quantized data set. It assigns the shortest code words to the most frequently occurring output values and thus reduces coding redundancy. The operation, of course, is reversible. Upon completion of the symbol coding step, the input image has been processed to remove each of the three redundancies.

Figure 3.2(a) shows the source encoding process as three successive operations, but all three operations are not necessarily included in every compression system. Recall, for example, that the quantizer must be omitted when error-free compression is desired. In addition, some compression techniques normally are modeled by merging blocks that are physically separate in

Fig. 3.2(a). In the predictive compression systems, for instance, the mapper and quantizer are often represented by a single block, which simultaneously performs both operations.

The source decoder shown in Fig. 3.2(b) contains only two components: a symbol decoder and an inverse mapper. These blocks perform, in reverse order, the inverse operations of the source encoder's symbol encoder and mapper blocks. Because quantization results in irreversible information loss, an inverse quantizer block is not included in the general source decoder model shown in Fig. 3.2(b).

The Channel Encoder and Decoder:

The channel encoder and decoder play an important role in the overall encoding-decoding process when the channel of Fig. 3.1 is noisy or prone to error. They are designed to reduce the impact of channel noise by inserting a controlled form of redundancy into the source encoded data. As the output of the source encoder contains little redundancy, it would be highly sensitive to transmission noise without the addition of this "controlled redundancy." One of the most useful channel encoding techniques was devised by R. W. Hamming (Hamming [1950]). It is based on appending enough bits to the data being encoded to ensure that some minimum number of bits must change between valid code words. Hamming showed, for example, that if 3 bits of redundancy are added to a 4-bit word, so that the distance between any two valid code words is 3, all single-bit errors can be detected and corrected. (By appending additional bits of redundancy, multiple-bit errors can be detected and corrected.) The 7-bit Hamming (7, 4) code word $h_1, h_2, h_3 \ldots, h_6, h_7$ associated with a 4-bit binary number $b_3b_2b_1b_0$ is

$$h_1 = b_3 \oplus b_2 \oplus b_0 \qquad h_3 = b_3$$
$$h_2 = b_3 \oplus b_1 \oplus b_0 \qquad h_5 = b_2$$
$$h_4 = b_2 \oplus b_1 \oplus b_0 \qquad h_6 = b_1$$
$$h_7 = b_0$$

where $\oplus$ denotes the exclusive OR operation. Note that bits h_1, h_2, and h_4 are even- parity bits for the bit fields $b_3 b_2 b_0$, $b_3b_1b_0$, and $b_2b_1b_0$, respectively. (Recall that a string of binary bits has even parity if the number of bits with a value of 1 is even.) To decode a Hamming encoded result, the channel decoder must check the encoded value for odd parity over the bit fields in which even parity was previously established. A single-bit error is indicated by a nonzero parity word $c_4c_2c_1$, where

$$c_1 = h_1 \oplus h_3 \oplus h_5 \oplus h_7$$
$$c_2 = h_2 \oplus h_3 \oplus h_6 \oplus h_7$$
$$c_4 = h_4 \oplus h_5 \oplus h_6 \oplus h_7.$$

If a nonzero value is found, the decoder simply complements the code word bit position indicated by the parity word. The decoded binary value is then extracted from the corrected code word as $h_3 h_5 h_6 h_7$.

4. Explain a method of generating variable length codes with an example.

Variable-Length Coding:

The simplest approach to error-free image compression is to reduce only coding redundancy. Coding redundancy normally is present in any natural binary encoding of the gray levels in an image. It can be eliminated by coding the gray levels. To do so requires construction of a variable-length code that assigns the shortest possible code words to the most probable gray levels. Here, we examine several optimal and near optimal techniques for constructing such a code. These techniques are formulated in the language of information theory. In practice, the source symbols may be either the gray levels of an image or the output of a gray-level mapping operation (pixel differences, run lengths, and so on).

Huffman coding:

The most popular technique for removing coding redundancy is due to Huffman (Huffman [1952]). When coding the symbols of an information source individually, Huffman coding yields the smallest possible number of code symbols per source symbol. In terms of the noiseless coding theorem, the resulting code is optimal for a fixed value of n, subject to the constraint that the source symbols be coded one at a time.

The first step in Huffman's approach is to create a series of source reductions by ordering the probabilities of the symbols under consideration and combining the lowest probability symbols into a single symbol that replaces them in the next source reduction. Figure 4.1 illustrates this process for binary coding (K-ary Huffman codes can also be constructed). At the far left, a hypothetical set of source symbols and their probabilities are ordered from top to bottom in terms of decreasing probability values. To form the first source reduction, the bottom two probabilities, 0.06 and 0.04, are combined to form a "compound symbol" with probability 0.1. This compound symbol and its associated probability are placed in the first source reduction column so that the

probabilities of the reduced source are also ordered from the most to the least probable. This process is then repeated until a reduced source with two symbols (at the far right) is reached.

             The second step in Huffman's procedure is to code each reduced source, starting with the smallest source and working back to the original source. The minimal length binary code for a two-symbol source, of course, is the symbols 0 and 1. As Fig. 4.2 shows, these symbols are assigned to the two symbols on the right (the assignment is arbitrary; reversing the order of the 0 and 1 would work just as well). As the reduced source symbol with probability 0.6 was generated by combining two symbols in the reduced source to its left, the 0 used to code it is now assigned to both of these symbols, and a 0 and 1 are arbitrarily

Fig.4.1 Huffman source reductions.

Fig.4.2 Huffman code assignment procedure.

appended to each to distinguish them from each other. This operation is then repeated for each reduced source until the original source is reached. The final code appears at the far left in Fig. 4.2. The average length of this code is

$$L_{avg} = (0.4)(1) + (0.3)(2) + (0.1)(3) + (0.1)(4) + (0.06)(5) + (0.04)(5)$$
$$= 2.2 \text{ bits/symbol}$$

and the entropy of the source is 2.14 bits/symbol. The resulting Huffman code efficiency is 0.973.

Huffman's procedure creates the optimal code for a set of symbols and probabilities subject to the constraint that the symbols be coded one at a time. After the code has been created, coding and/or decoding is accomplished in a simple lookup table manner. The code itself is an instantaneous uniquely decodable block code. It is called a block code because each source symbol is mapped into a fixed sequence of code symbols. It is instantaneous, because each code word in a string of code symbols can be decoded without referencing succeeding symbols. It is uniquely decodable, because any string of code symbols can be decoded in only one way. Thus, any string of Huffman encoded symbols can be decoded by examining the individual symbols of the string in a left to right manner. For the binary code of Fig. 4.2, a left-to-right scan of the encoded string 010100111100 reveals that the first valid code word is 01010, which is the code for symbol a_3 .The next valid code is 011, which corresponds to symbol a_1. Continuing in this manner reveals the completely decoded message to be $a_3a_1a_2a_2a_6$.

5. Explain arithmetic encoding process with an example.

Arithmetic coding:

Unlike the variable-length codes described previously, arithmetic coding generates nonblock codes. In arithmetic coding, which can be traced to the work of Elias, a one-to-one correspondence between source symbols and code words does not exist. Instead, an entire sequence of source symbols (or message) is assigned a single arithmetic code word. The code word itself defines an interval of real numbers between 0 and 1. As the number of symbols in the message increases, the interval used to represent it becomes smaller and the number of information units (say, bits) required to represent the interval becomes larger. Each symbol of the message reduces the size of the interval in accordance with its probability of occurrence. Because the technique does not require, as does Huffman's approach, that each source symbol translate into an integral number of code symbols (that is, that the symbols be coded one at a time), it achieves (but only in theory) the bound established by the noiseless coding theorem.
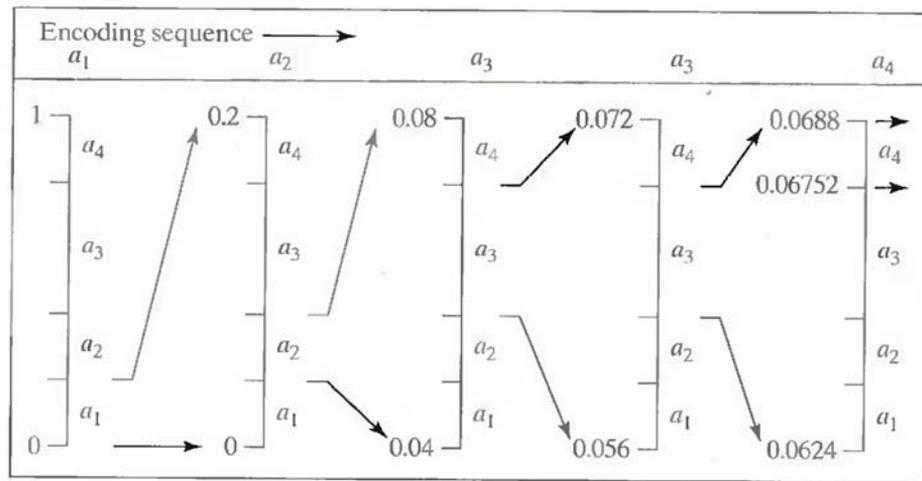
Fig.5.1 Arithmetic coding procedure

Figure 5.1 illustrates the basic arithmetic coding process. Here, a five-symbol sequence or message, $a_1a_2a_3a_3a_4$, from a four-symbol source is coded. At the start of the coding process, the message is assumed to occupy the entire half-open interval [0, 1). As Table 5.2 shows, this interval is initially subdivided into four regions based on the probabilities of each source symbol. Symbol ax, for example, is associated with subinterval [0, 0.2). Because it is the first symbol of the message being coded, the message interval is initially narrowed to [0, 0.2). Thus in Fig. 5.1 [0, 0.2) is expanded to the full height of the figure and its end points labeled by the values of the narrowed range. The narrowed range is then subdivided in accordance with the original source symbol probabilities and the process continues with the next message symbol.

Source Symbol	Probability	Initial Subinterval
a_1	0.2	[0.0, 0.2)
a_2	0.2	[0.2, 0.4)
a_3	0.4	[0.4, 0.8)
a_4	0.2	[0.8, 1.0)

Table 5.1 Arithmetic coding example

In this manner, symbol a_2 narrows the subinterval to [0.04, 0.08), a_3 further narrows it to [0.056, 0.072), and so on. The final message symbol, which must be reserved as a special end-of-

message indicator, narrows the range to [0.06752, 0.0688). Of course, any number within this subinterval—for example, 0.068—can be used to represent the message.

In the arithmetically coded message of Fig. 5.1, three decimal digits are used to represent the five-symbol message. This translates into 3/5 or 0.6 decimal digits per source symbol and compares favorably with the entropy of the source, which is 0.58 decimal digits or 10-ary units/symbol. As the length of the sequence being coded increases, the resulting arithmetic code approaches the bound established by the noiseless coding theorem.

In practice, two factors cause coding performance to fall short of the bound: (1) the addition of the end-of-message indicator that is needed to separate one message from another; and (2) the use of finite precision arithmetic. Practical implementations of arithmetic coding address the latter problem by introducing a scaling strategy and a rounding strategy (Langdon and Rissanen [1981]). The scaling strategy renormalizes each subinterval to the [0, 1) range before subdividing it in accordance with the symbol probabilities. The rounding strategy guarantees that the truncations associated with finite precision arithmetic do not prevent the coding subintervals from being represented accurately.

6. Explain LZW coding with an example.

LZW Coding:

The technique, called Lempel-Ziv-Welch (LZW) coding, assigns fixed-length code words to variable length sequences of source symbols but requires no a priori knowledge of the probability of occurrence of the symbols to be encoded. LZW compression has been integrated into a variety of mainstream imaging file formats, including the graphic interchange format (GIF), tagged image file format (TIFF), and the portable document format (PDF).

LZW coding is conceptually very simple (Welch [1984]). At the onset of the coding process, a codebook or "dictionary" containing the source symbols to be coded is constructed. For 8-bit monochrome images, the first 256 words of the dictionary are assigned to the gray values 0, 1, 2..., and 255. As the encoder sequentially examines the image's pixels, gray-level sequences that are not in the dictionary are placed in algorithmically determined (e.g., the next unused) locations. If the first two pixels of the image are white, for instance, sequence "255-255" might be assigned to location 256, the address following the locations reserved for gray levels 0 through 255. The next time that two consecutive white pixels are encountered, code word 256, the address of the location containing sequence 255-255, is used to represent them. If a 9-bit, 512-word dictionary is employed in the coding process, the original (8 + 8) bits that were used to represent the two pixels are replaced by a single 9-bit code word. Cleary, the size of the

dictionary is an important system parameter. If it is too small, the detection of matching gray-level sequences will be less likely; if it is too large, the size of the code words will adversely affect compression performance.

Consider the following 4 x 4, 8-bit image of a vertical edge:

39	39	126	126
39	39	126	126
39	39	126	126
39	39	126	126

Table 6.1 details the steps involved in coding its 16 pixels. A 512-word dictionary with the following starting content is assumed:

Dictionary Location	Entry
0	0
1	1
⋮	⋮
255	255
256	—
⋮	⋮
511	—

Locations 256 through 511 are initially unused. The image is encoded by processing its pixels in a left-to-right, top-to-bottom manner. Each successive gray-level value is concatenated with a variable—column 1 of Table 6.1 —called the "currently recognized sequence." As can be seen, this variable is initially null or empty. The dictionary is searched for each concatenated sequence and if found, as was the case in the first row of the table, is replaced by the newly concatenated and recognized (i.e., located in the dictionary) sequence. This was done in column 1 of row 2.

Currently Recognized Sequence	Pixel Being Processed	Encoded Output	Dictionary Location (Code Word)	Dictionary Entry
	39			
39	39	39	256	39-39
39	126	39	257	39-126
126	126	126	258	126-126
126	39	126	259	126-39
39	39			
39-39	126	256	260	39-39-126
126	126			
126-126	39	258	261	126-126-39
39	39			
39-39	126			
39-39-126	126	260	262	39-39-126-126
126	39			
126-39	39	259	263	126-39-39
39	126			
39-126	126	257	264	39-126-126
126		126		

Table 6.1 LZW coding example

No output codes are generated, nor is the dictionary altered. If the concatenated sequence is not found, however, the address of the currently recognized sequence is output as the next encoded value, the concatenated but unrecognized sequence is added to the dictionary, and the currently recognized sequence is initialized to the current pixel value. This occurred in row 2 of the table. The last two columns detail the gray-level sequences that are added to the dictionary when scanning the entire 4 x 4 image. Nine additional code words are defined. At the conclusion of coding, the dictionary contains 265 code words and the LZW algorithm has successfully identified several repeating gray-level sequences—leveraging them to reduce the original 128-bit image lo 90 bits (i.e., 10 9-bit codes). The encoded output is obtained by reading the third column from top to bottom. The resulting compression ratio is 1.42:1.

A unique feature of the LZW coding just demonstrated is that the coding dictionary or code book is created while the data are being encoded. Remarkably, an LZW decoder builds an identical decompression dictionary as it decodes simultaneously the encoded data stream. . Although not needed in this example, most practical applications require a strategy for handling dictionary overflow. A simple solution is to flush or reinitialize the dictionary when it becomes full and continue coding with a new initialized dictionary. A more complex option is

to monitor compression performance and flush the dictionary when it becomes poor or unacceptable. Alternately, the least used dictionary entries can be tracked and replaced when necessary.

7. Explain the concept of bit plane coding method.

Bit-Plane Coding:

An effective technique for reducing an image's interpixel redundancies is to process the image's bit planes individually. The technique, called bit-plane coding, is based on the concept of decomposing a multilevel (monochrome or color) image into a series of binary images and compressing each binary image via one of several well-known binary compression methods.

Bit-plane decomposition:

The gray levels of an m-bit gray-scale image can be represented in the form of the base 2 polynomial

$$a_{m-1}2^{m-1} + a_{m-2}2^{m-2} + \cdots + a_1 2^1 + a_0 2^0.$$

Based on this property, a simple method of decomposing the image into a collection of binary images is to separate the m coefficients of the polynomial into m 1-bit bit planes. The zeroth-order bit plane is generated by collecting the a_0 bits of each pixel, while the (m - 1) st-order bit plane contains the a_{m-1}, bits or coefficients. In general, each bit plane is numbered from 0 to m-1 and is constructed by setting its pixels equal to the values of the appropriate bits or polynomial coefficients from each pixel in the original image. The inherent disadvantage of this approach is that small changes in gray level can have a significant impact on the complexity of the bit planes. If a pixel of intensity 127 (01111111) is adjacent to a pixel of intensity 128 (10000000), for instance, every bit plane will contain a corresponding 0 to 1 (or 1 to 0) transition. For example, as the most significant bits of the two binary codes for 127 and 128 are different, bit plane 7 will contain a zero-valued pixel next to a pixel of value 1, creating a 0 to 1 (or 1 to 0) transition at that point.

An alternative decomposition approach (which reduces the effect of small gray-level variations) is to first represent the image by an m-bit Gray code. The m-bit Gray code $g_{m-1}...g_2g_1g_0$ that corresponds to the polynomial in Eq. above can be computed from

$$g_i = a_i \oplus a_{i+1} \qquad 0 \le i \le m - 2$$
$$g_{m-1} = a_{m-1}.$$

Here, $\oplus$ denotes the exclusive OR operation. This code has the unique property that successive code words differ in only one bit position. Thus, small changes in gray level are less likely to affect all m bit planes. For instance, when gray levels 127 and 128 are adjacent, only the 7th bit plane will contain a 0 to 1 transition, because the Gray codes that correspond to 127 and 128 are 11000000 and 01000000, respectively.

8. Explain about lossless predictive coding.

Lossless Predictive Coding:

The error-free compression approach does not require decomposition of an image into a collection of bit planes. The approach, commonly referred to as lossless predictive coding, is based on eliminating the interpixel redundancies of closely spaced pixels by extracting and coding only the new information in each pixel. The new information of a pixel is defined as the difference between the actual and predicted value of that pixel.

Figure 8.1 shows the basic components of a lossless predictive coding system. The system consists of an encoder and a decoder, each containing an identical predictor. As each successive pixel of the input image, denoted f_n, is introduced to the encoder, the predictor generates the anticipated value of that pixel based on some number of past inputs. The output of the predictor is then rounded to the nearest integer, denoted $f\hat{}_n$ and used to form the difference or prediction error which is coded using a variable-length code (by the symbol encoder) to generate the next element of the compressed data stream.

$$e_n = f_n - \hat{f}_n,$$
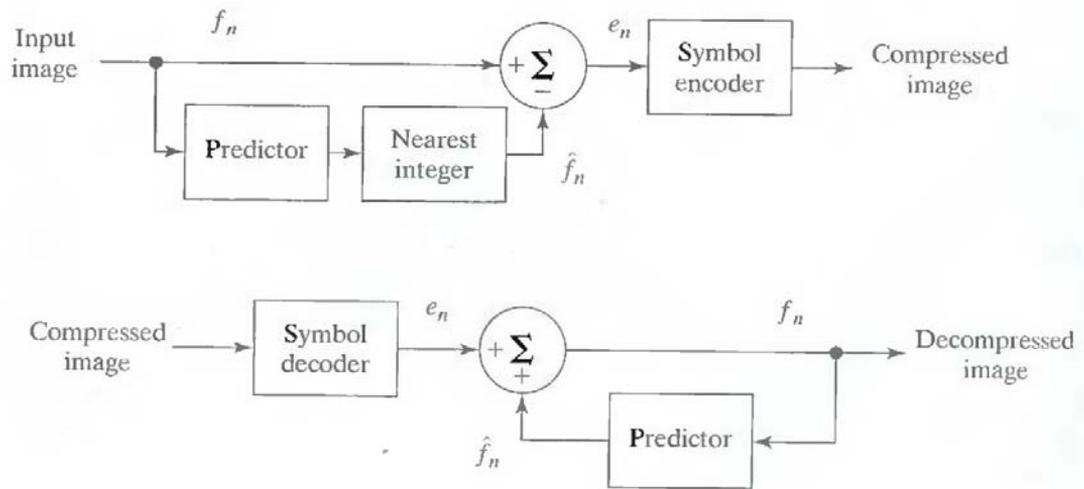
Fig.8.1 A lossless predictive coding model: (a) encoder; (b) decoder

The decoder of Fig. 8.1 (b) reconstructs en from the received variable-length code words and performs the inverse operation

$$f_n = e_n + \hat{f}_n.$$

Various local, global, and adaptive methods can be used to generate f^n. In most cases, however, the prediction is formed by a linear combination of m previous pixels. That is,

$$\hat{f}_n = \text{round}\left[\sum_{i=1}^{m} \alpha_i f_{n-i}\right]$$

where m is the order of the linear predictor, round is a function used to denote the rounding or nearest integer operation, and the α_i, for i = 1,2,..., m are prediction coefficients. In raster scan applications, the subscript n indexes the predictor outputs in accordance with their time of occurrence. That is, f_n, f^n and e_n in Eqns. above could be replaced with the more explicit notation f (t), f^(t), and e (t), where t represents time. In other cases, n is used as an index on the spatial coordinates and/or frame number (in a time sequence of images) of an image. In 1-D linear predictive coding, for example, Eq. above can be written as

$$\hat{f}_n(x, y) = \text{round}\left[\sum_{i=1}^{m} \alpha_i f(x, y - i)\right]$$

where each subscripted variable is now expressed explicitly as a function of spatial coordinates x and y. The Eq. indicates that the 1-D linear prediction f(x, y) is a function of the previous pixels on the current line alone. In 2-D predictive coding, the prediction is a function of the previous pixels in a left-to-right, top-to-bottom scan of an image. In the 3-D case, it is based on these pixels and the previous pixels of preceding frames. Equation above cannot be evaluated for the first m pixels of each line, so these pixels must be coded by using other means (such as a Huffman code) and considered as an overhead of the predictive coding process. A similar comment applies to the higher-dimensional cases.

9. Explain about lossy predictive coding.

Lossy Predictive Coding:

In this type of coding, we add a quantizer to the lossless predictive model and examine the resulting trade-off between reconstruction accuracy and compression performance. As Fig.9 shows, the quantizer, which absorbs the nearest integer function of the error-free encoder, is inserted between the symbol encoder and the point at which the prediction error is formed. It maps the prediction error into a limited range of outputs, denoted $e\hat{}_n$ which establish the amount of compression and distortion associated with lossy predictive coding.
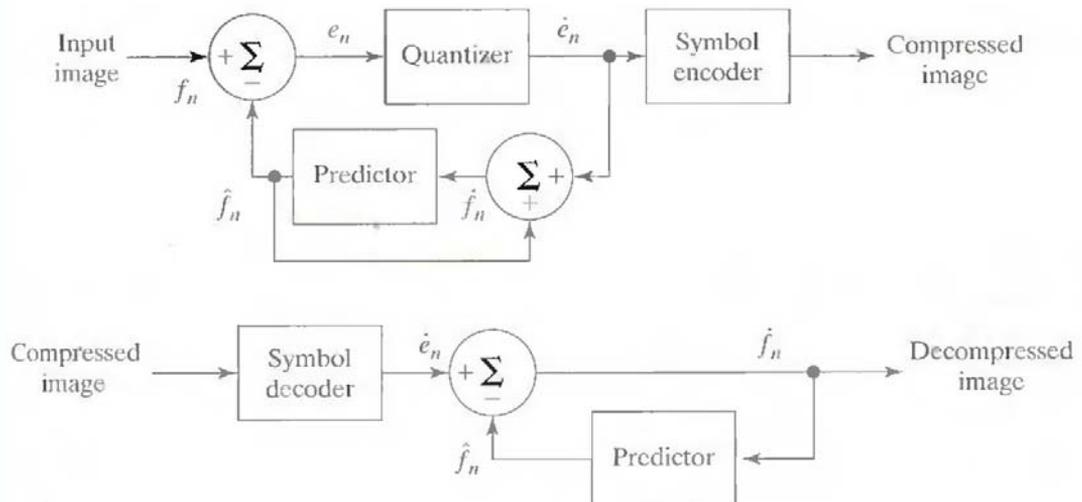
Fig. 9 A lossy predictive coding model: (a) encoder and (b) decoder.

In order to accommodate the insertion of the quantization step, the error-free encoder of figure must be altered so that the predictions generated by the encoder and decoder are equivalent. As Fig.9 (a) shows, this is accomplished by placing the lossy encoder's predictor within a feedback loop, where its input, denoted $\dot{f}_n$, is generated as a function of past predictions and the corresponding quantized errors. That is,

$$\dot{f}_n = \dot{e}_n + \hat{f}_n$$

 This closed loop configuration prevents error buildup at the decoder's output. Note from Fig. 9 (b) that the output of the decoder also is given by the above Eqn.

Optimal predictors:

The optimal predictor used in most predictive coding applications minimizes the encoder's mean-square prediction error

$$E\{e_n^2\} = E\{[f_n - \hat{f}_n]^2\}$$

subject to the constraint that

$$\dot{f}_n = \dot{e}_n + \hat{f}_n \approx e_n + \hat{f}_n = f_n$$

and

$$\hat{f}_n = \sum_{i=1}^{m} \alpha_i f_{n-i}.$$

That is, the optimization criterion is chosen to minimize the mean-square prediction error, the quantization error is assumed to be negligible ($\dot{e}_n \approx e_n$), and the prediction is constrained to a linear combination of m previous pixels.1 These restrictions are not essential, but they simplify the analysis considerably and, at the same time, decrease the computational complexity of the predictor. The resulting predictive coding approach is referred to as differential pulse code modulation (DPCM).

10. Explain with a block diagram about transform coding system.

Transform Coding:

All the predictive coding techniques operate directly on the pixels of an image and thus are spatial domain methods. In this coding, we consider compression techniques that are based on modifying the transform of an image. In transform coding, a reversible, linear transform (such as the Fourier transform) is used to map the image into a set of transform coefficients, which are then quantized and coded. For most natural images, a significant number of the coefficients have small magnitudes and can be coarsely quantized (or discarded entirely) with little image distortion. A variety of transformations, including the discrete Fourier transform (DFT), can be used to transform the image data.
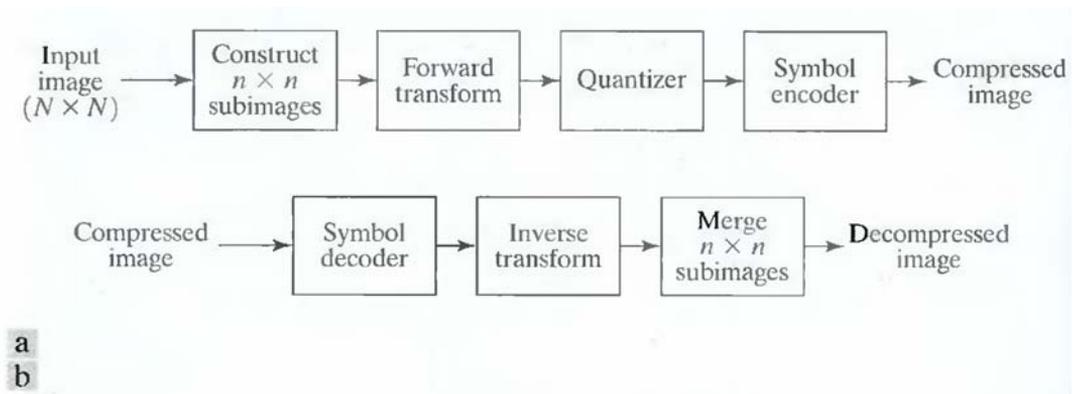
Fig. 10 A transform coding system: (a) encoder; (b) decoder.

Figure 10 shows a typical transform coding system. The decoder implements the inverse sequence of steps (with the exception of the quantization function) of the encoder, which performs four relatively straightforward operations: subimage decomposition, transformation, quantization, and coding. An N X N input image first is subdivided into subimages of size n X n, which are then transformed to generate $(N/n)^2$ subimage transform arrays, each of size n X n. The goal of the transformation process is to decorrelate the pixels of each subimage, or to pack as much information as possible into the smallest number of transform coefficients. The quantization stage then selectively eliminates or more coarsely quantizes the coefficients that carry the least information. These coefficients have the smallest impact on reconstructed subimage quality. The encoding process terminates by coding (normally using a variable-length code) the quantized coefficients. Any or all of the transform encoding steps can be adapted to

local image content, called adaptive transform coding, or fixed for all subimages, called nonadaptive transform coding.

11. Explain about wavelet coding.

Wavelet Coding:

The wavelet coding is based on the idea that the coefficients of a transform that decorrelates the pixels of an image can be coded more efficiently than the original pixels themselves. If the transform's basis functions—in this case wavelets—pack most of the important visual information into a small number of coefficients, the remaining coefficients can be quantized coarsely or truncated to zero with little image distortion.

Figure 11 shows a typical wavelet coding system. To encode a 2^J X 2^J image, an analyzing wavelet, Ψ, and minimum decomposition level, J - P, are selected and used to compute the image's discrete wavelet transform. If the wavelet has a complimentary scaling function φ, the fast wavelet transform can be used. In either case, the computed transform converts a large portion of the original image to horizontal, vertical, and diagonal decomposition coefficients with zero mean and Laplacian-like distributions.
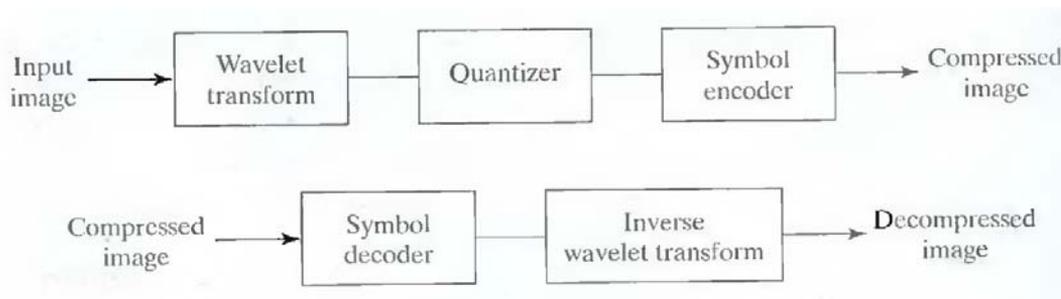
Fig.11 A wavelet coding system: (a) encoder; (b) decoder.

 Since many of the computed coefficients carry little visual information, they can be quantized and coded to minimize intercoefficient and coding redundancy. Moreover, the quantization can be adapted to exploit any positional correlation across the P decomposition levels. One or more of the lossless coding methods, including run-length, Huffman, arithmetic, and bit-plane coding, can be incorporated into the final symbol coding step. Decoding is accomplished by inverting the encoding operations—with the exception of quantization, which cannot be reversed exactly.

                        The principal difference between the wavelet-based system and the transform coding system is the omission of the transform coder's subimage processing stages.

Because wavelet transforms are both computationally efficient and inherently local (i.e., their basis functions are limited in duration), subdivision of the original image is unnecessary.
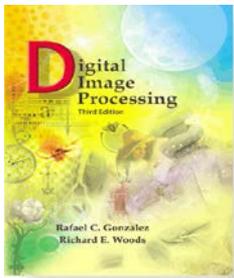
The term morphology refers to the description of the properties of shape and structure of any objects. Operations of mathematical morphology were originally defined as operations on sets, but it soon became clear that they are also useful in the processing tasks of the set of points in the two-dimensional space. Sets in mathematical morphology represent objects in the image. It is easy to see that the set of all background pixels of binary image is one of the options for a full description.

Mathematical morphology is used to extract some properties of the image, useful for its presentation and descriptions. For example, contours, skeletons and convex hulls. Also morphological methods are used in the preliminary and final image processing. For example, morphological filtering, thickening or thinning.
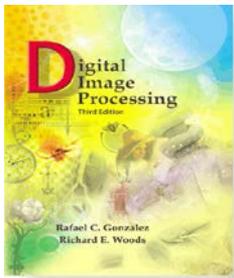
Reflection and translation

$$\hat{B} = \{w \mid w = -b,\, for\ b \in B\}$$

$$(B)_z = \{c \mid c = b + z,\, for\ b \in B\}$$
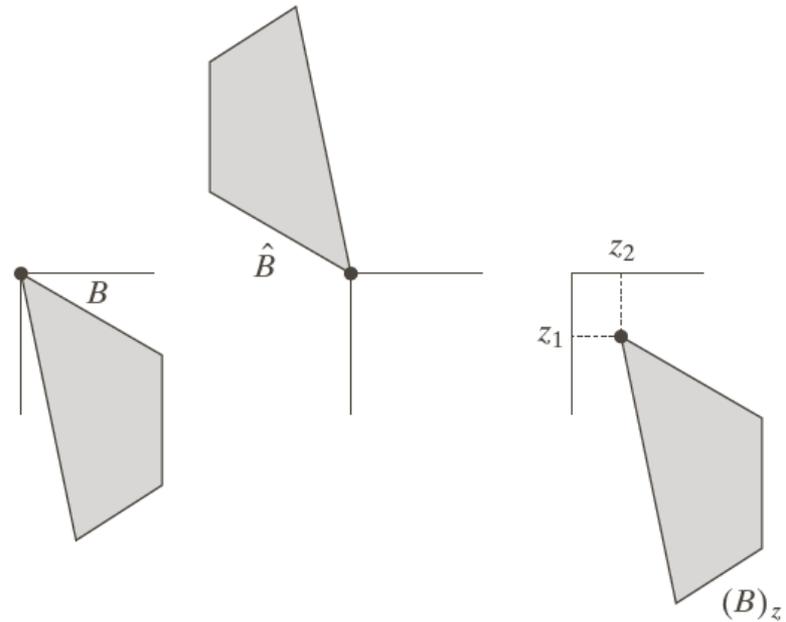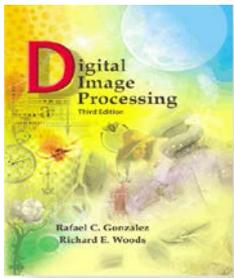
Basic Operations: union, intersection, complement, difference

Input Images
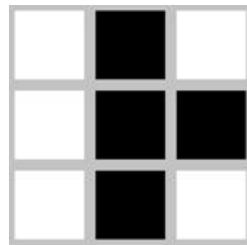
A   B
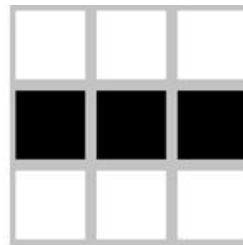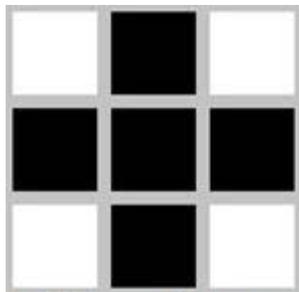
C = A ∪ B   C = A ∩ B   C = A^c C = A\B
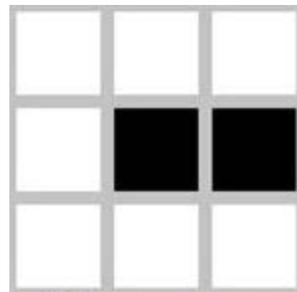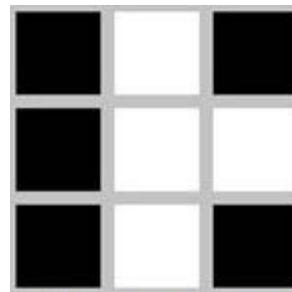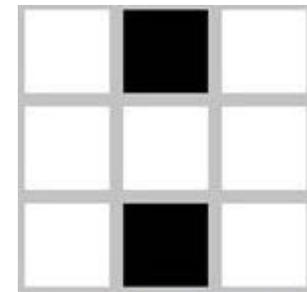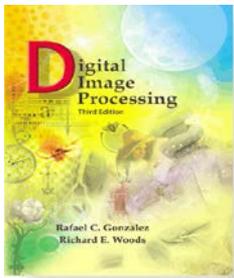
Structural Element

The input data for the mathematical morphology are the two images: the processed image and a **structural element.**
Typically, a structural element is much smaller than the processed. The shape of the structured element can be arbitrary, as long as it can be represented as a binary image of a given size.
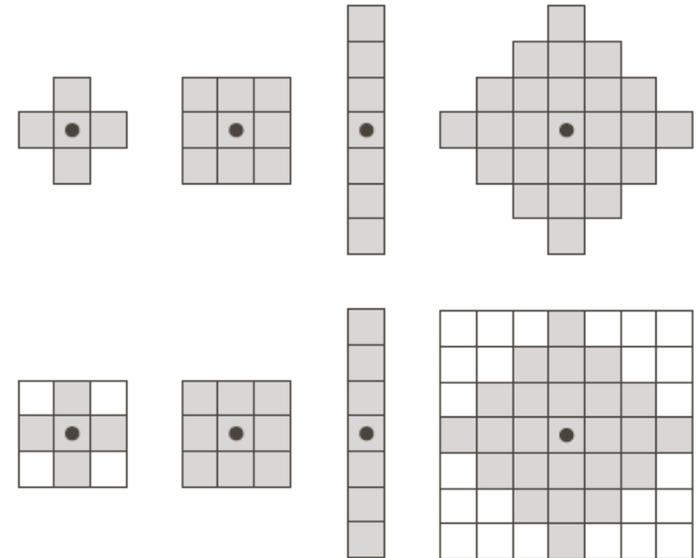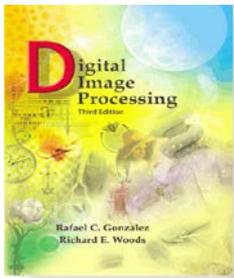
The structural element, S, is applied to all pixels of binary image. Every time the origin of the structural element is combined with a single binary pixel, the entire structural element is wrapped and subsequent alteration of the corresponding pixels of binary image. The results of logical addition is written into the output binary image, which was originally initialized to zero.
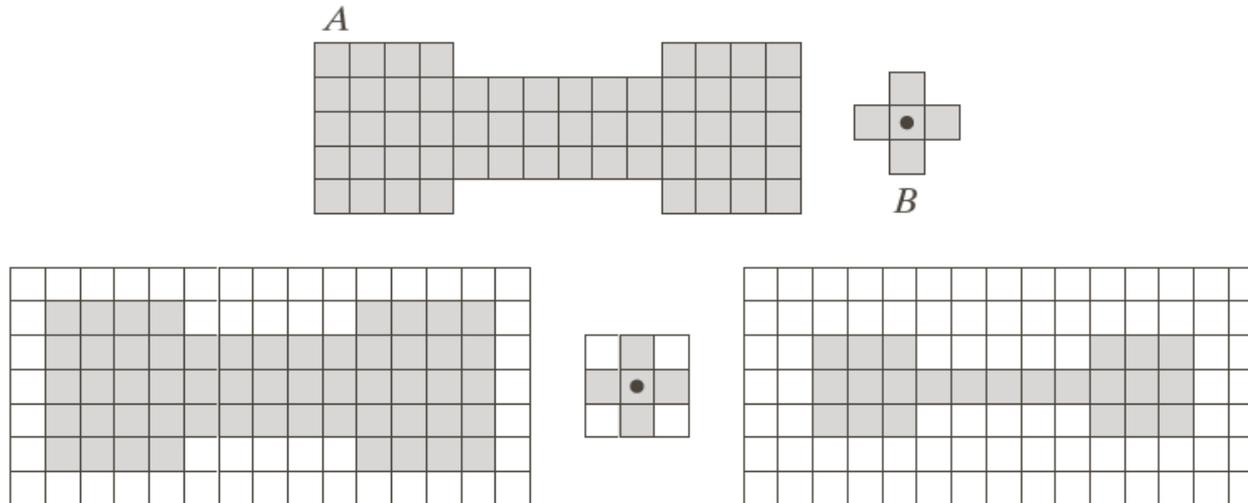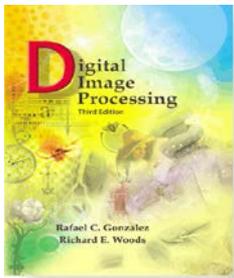
Erosion

For erosion, the structural element passes through all pixels of the image. The pixel corresponds to the reference of the structure element is turned <u>on</u> if the entire structure element falls with foreground area (pixels).

$$A \ominus B = \{z \,|\, (B)_z \subseteq A\}$$

$$A \ominus B = \{z \,|\, (B)_z \cap A^c = \emptyset\}$$

Erosion Example:

Input	1	0	0	1	1	1	1	0	1	1

SE | 1 | 1 | 1 |

Output		0								

Input	1	0	0	1	1	1	1	0	1	1

SE | 1 | 1 | 1 |

Output		0	0							

Input	1	0	0	1	1	1	1	0	1	1

SE | 1 | 1 | 1 |

Output		0	0	0						

Input	1	0	0	1	1	1	1	0	1	1

SE | 1 | 1 | 1 |

Output		0	0	0	1					

Erosion Example:

Panel 1:

Input	1	0	0	1	1	1	1	0	1	1

SE	1	1	1

Output		0	0	0	1	1				

Panel 2:

Input	1	0	0	1	1	1	1	0	1	1

SE	1	1	1

Output		0	0	0	1	1	0			

Panel 3:

Input	1	0	0	1	1	1	1	0	1	1

SE	1	1	1

Output		0	0	0	1	1	0	0		

Panel 4:

Input	1	0	0	1	1	1	1	0	1	1

SE	1	1	1

Output		0	0	0	1	1	0	0	0	
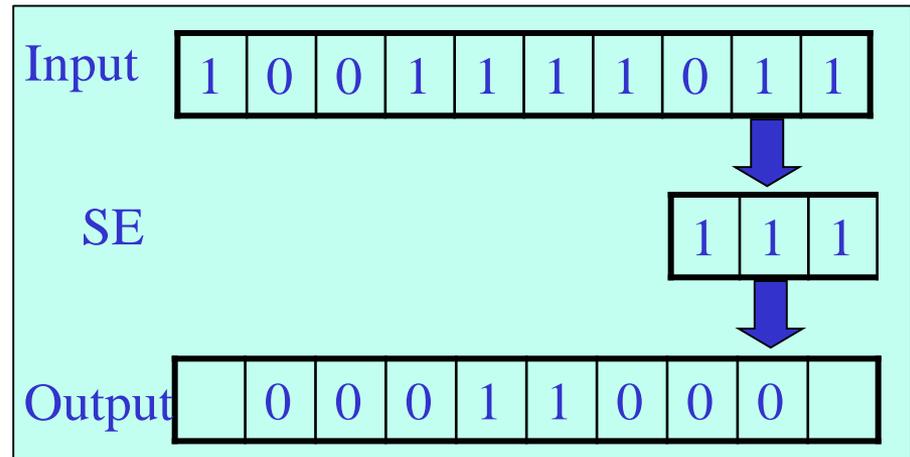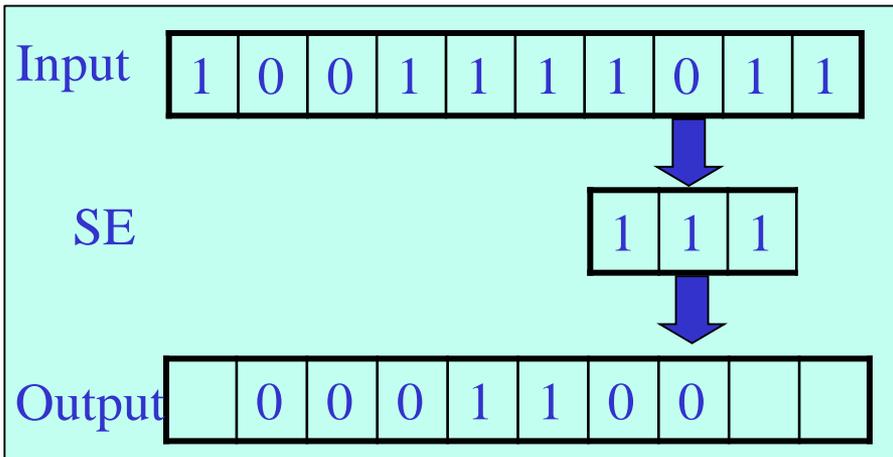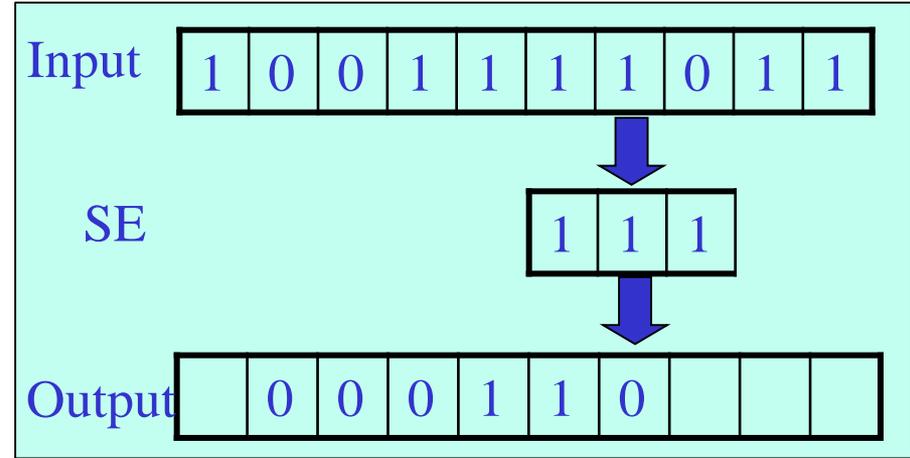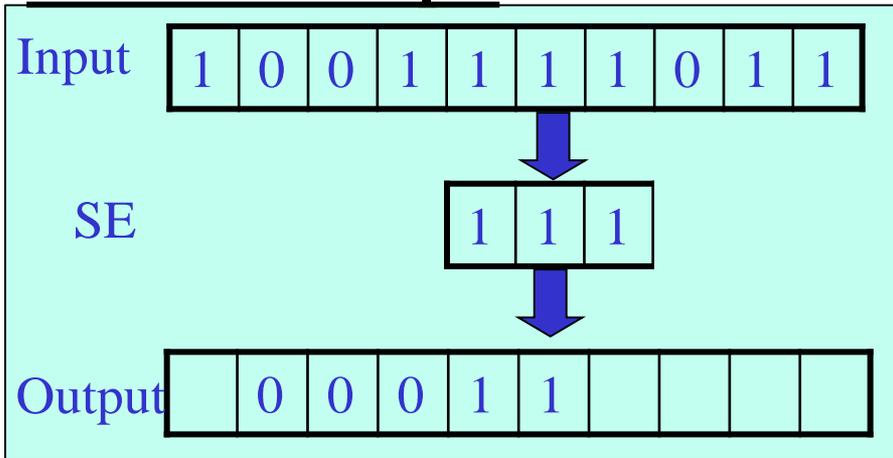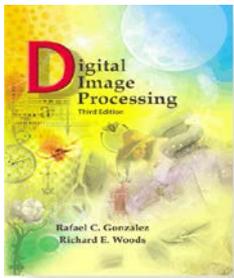
Erosion examples

Erosion examples

Image (486×486) eroded using squared structured element of size 11×11, 15×15, 45×45,
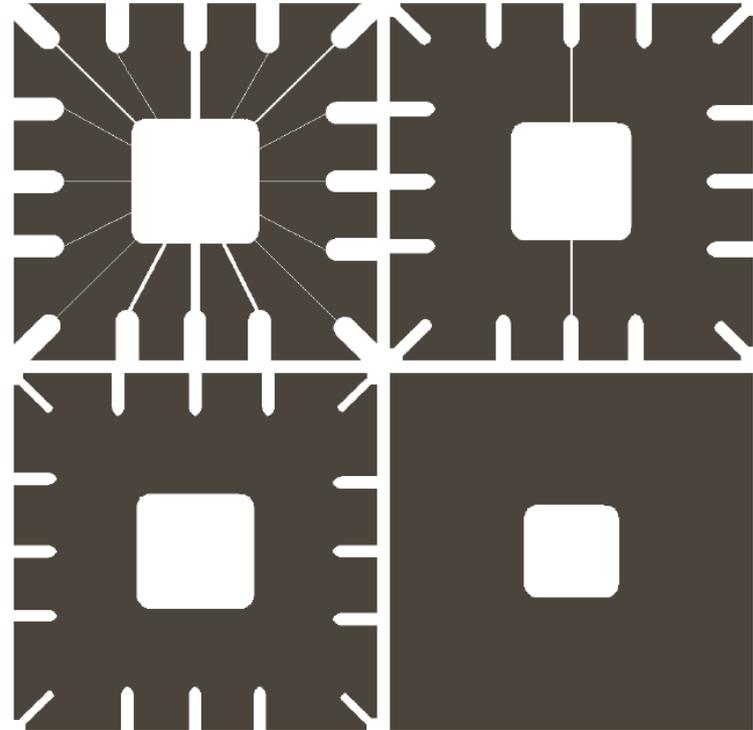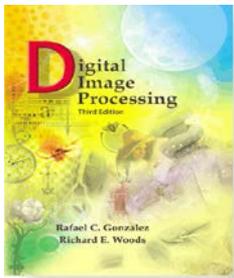
Dilation

The dilation of A and B is set of all displacement z, such that $\hat{B}$ and A overlap by at least one element.

$$A \oplus B = \{z \mid (\hat{B})_z \cap A \neq \emptyset\}$$

$$A \oplus B = \{z \mid [(\hat{B})_z \cap A] \subseteq A\}$$

Dilation Example:

Quadrant 1 (top-left):

Input: 1 0 0 0 1 1 1 0 1 1

SE: 1 1 1

Output: _ 1 _ _ _ _ _ _ _ _

Quadrant 2 (top-right):

Input: 1 0 0 0 1 1 1 0 1 1

SE: 1 1 1

Output: _ 1 0 _ _ _ _ _ _ _

Quadrant 3 (bottom-left):

Input: 1 0 0 0 1 1 1 0 1 1

SE: 1 1 1

Output: _ 1 0 1 _ _ _ _ _ _

Quadrant 4 (bottom-right):

Input: 1 0 0 0 1 1 1 0 1 1

SE: 1 1 1

Output: _ 1 0 1 1 _ _ _ _ _

<u>Dilation Example:</u>

Panel 1:

Input	1	0	0	0	1	1	1	0	1	1

SE	1	1	1

Output		1	0	1	1	1				

Panel 2:

Input	1	0	0	0	1	1	1	0	1	1

SE	1	1	1

Output		1	0	1	1	1	1			

Panel 3:

Input	1	0	0	0	1	1	1	0	1	1

SE	1	1	1

Output		1	0	1	1	1	1	1		

Panel 4:

Input	1	0	0	0	1	1	1	0	1	1

SE	1	1	1

Output		1	0	1	1	1	1	1	1	
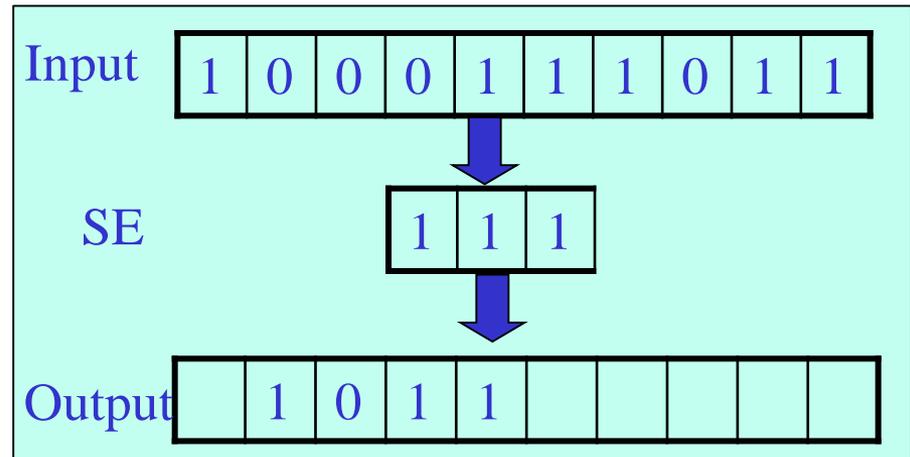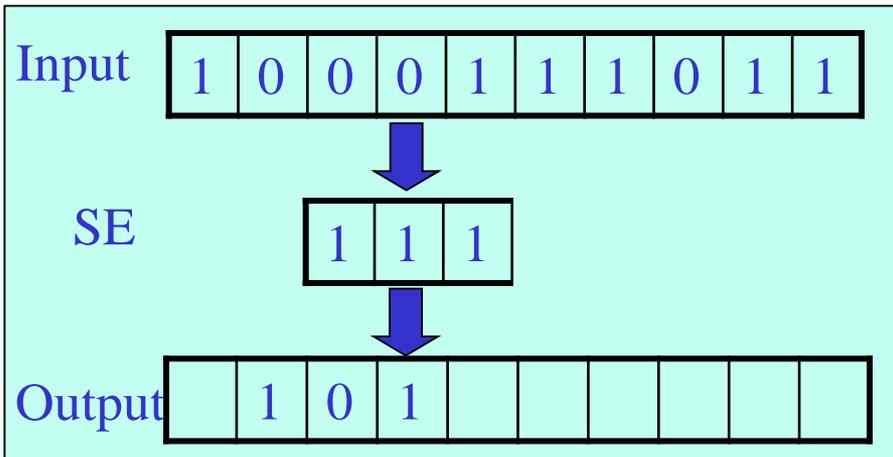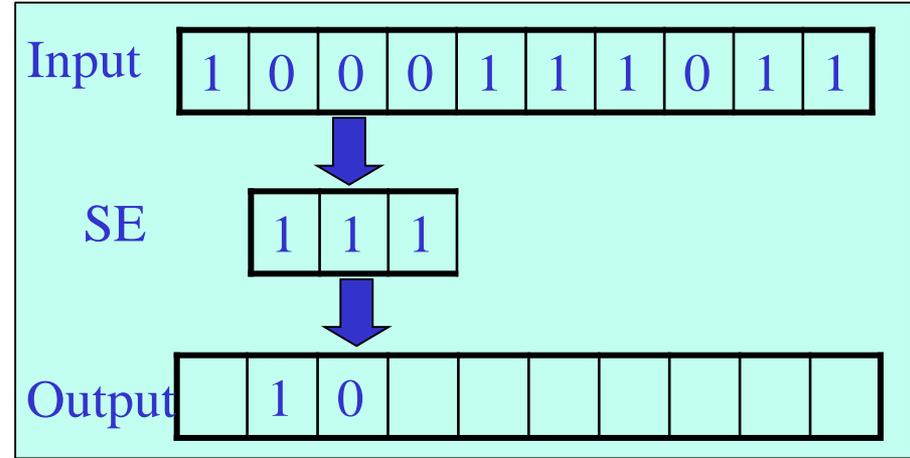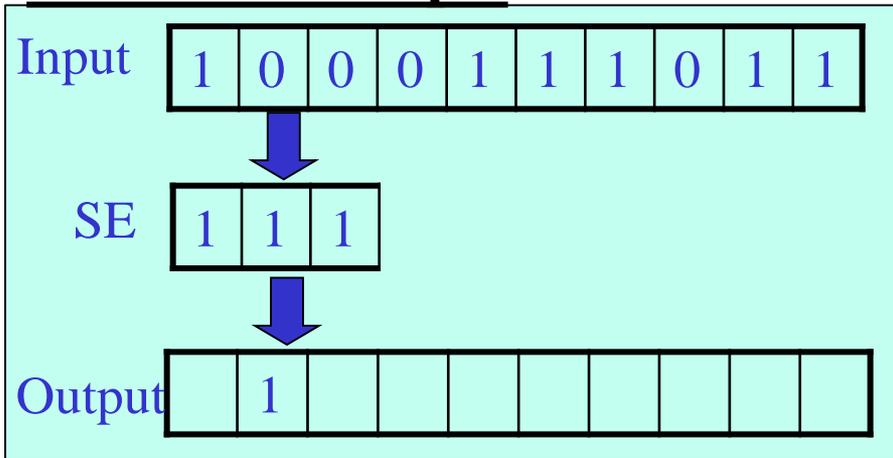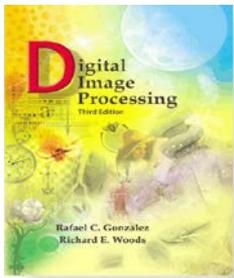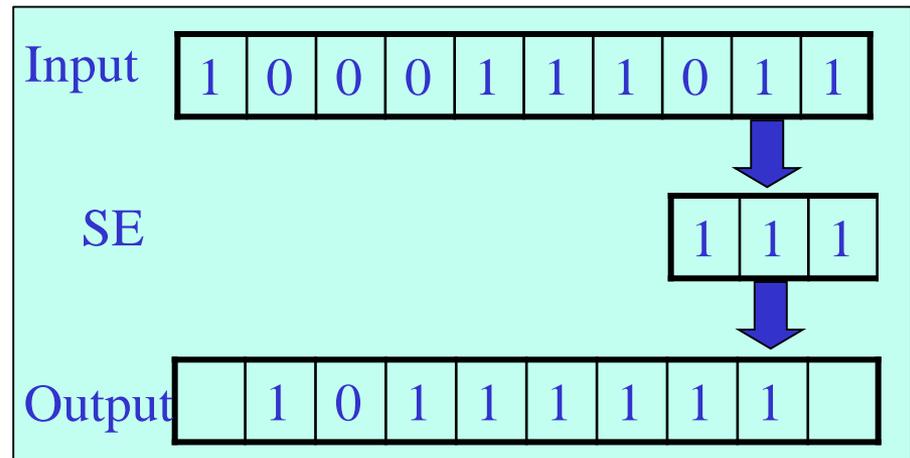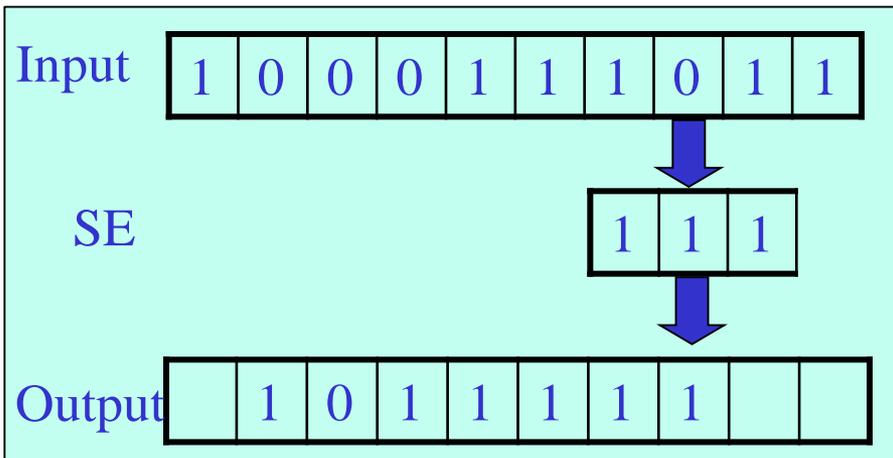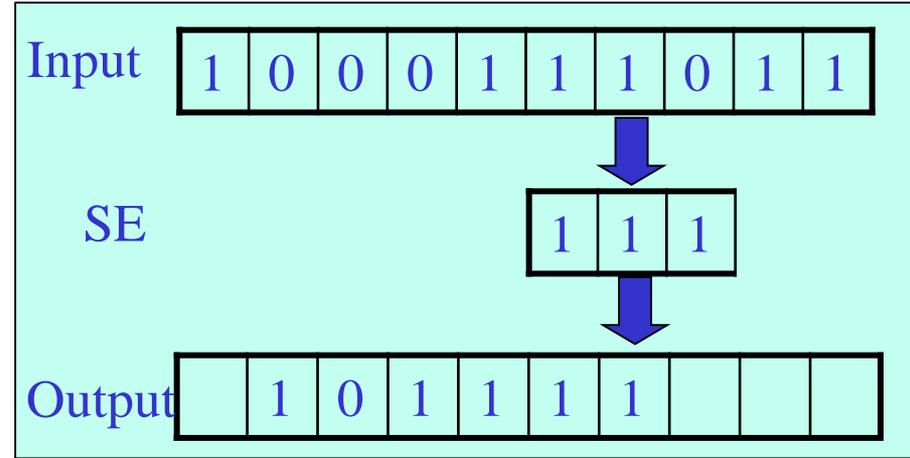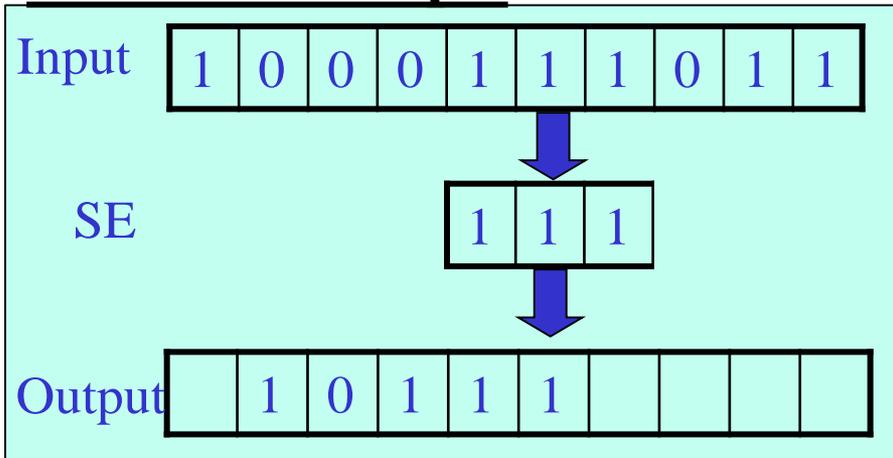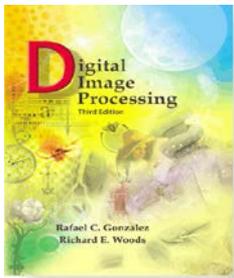
<u>Dilation Example:</u>

Input Image   SE   Output Image

Opening:

Opening aims to spot and noise, but in less destructive manner

The opening of *A* by *B* is obtained by the erosion of *A* by *B*, followed by dilation of the resulting image by *B*:

Opening is performed by sliding the structuring element (SE) around *inside* each foreground region.

- The pixels that can be covered by the SE, while SE being entirely within the foreground region will be preserved.
- The foreground pixels which **can *not* be reached** by the SE without lapping over the boundary of the foreground object will be eroded away.

Note that applying opening again will not change the image

Opening Example (rolling inside):

$$A \circ B = \cup\{(B)_z | (B)_z \subseteq A\}$$

Translates of *B* in *A*

boundary   Final Results

Opening Example (3×3):

Opening Example (11 pixel dimeter):

Opening Example (11 pixel diameter):

3×9

9×3

Closing:

Closing aims to remove holes, it tends to enlarge regions and shrink background

The closing of *A* by *B* is obtained by the dilation of *A* by *B*, followed by erosion of the resulting structure by *B*:

Closing performed by sliding the structuring element (SE) around *outside* each foreground region.

- Background pixels that can be covered by the SE, while SE being entirely within the background region will be preserved.
- The background pixels which **can *not* be reached** by the SE without lapping over the boundary of the foreground object will be turned to foreground.

Note that applying clsoing again will not change the image
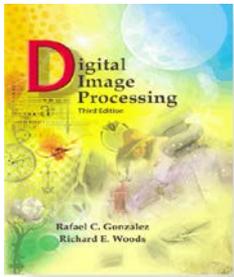
Closing Example (rolling outside):
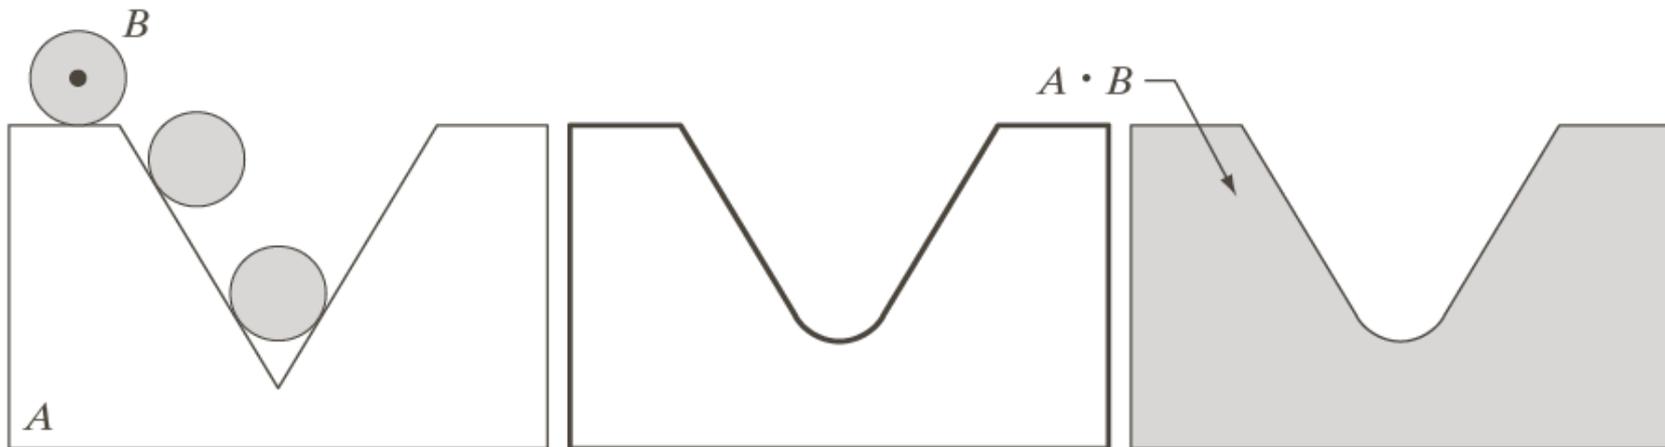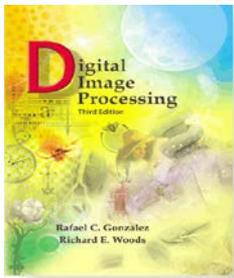
Closing Example (SE: 3×3 squre):

Closing Example (SE: 22 pixel disc):
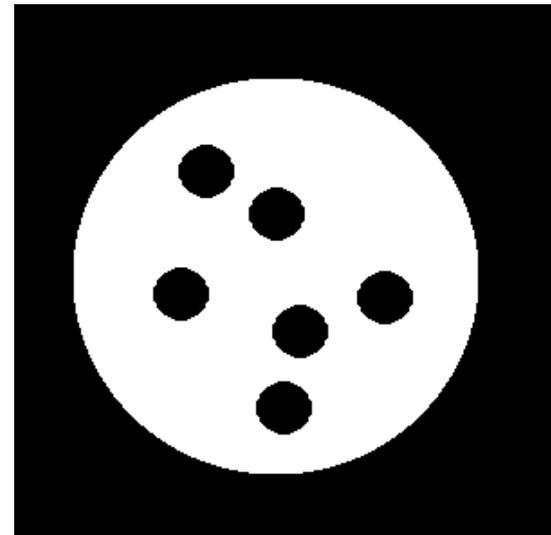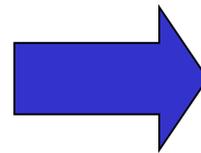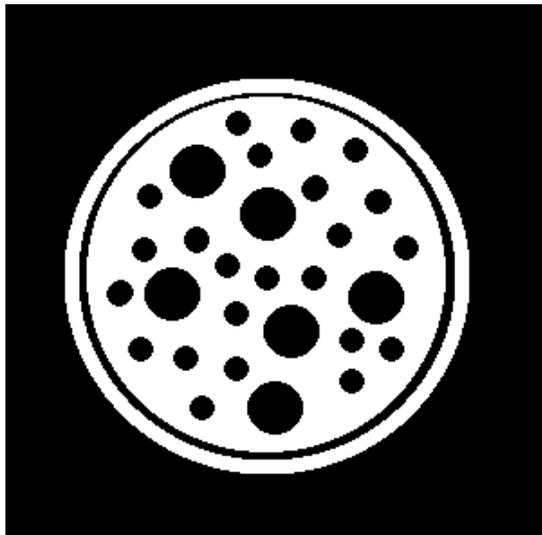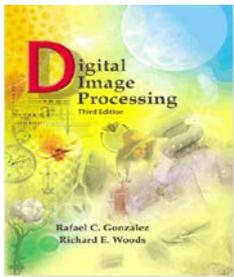
Opening and Closing Examples:
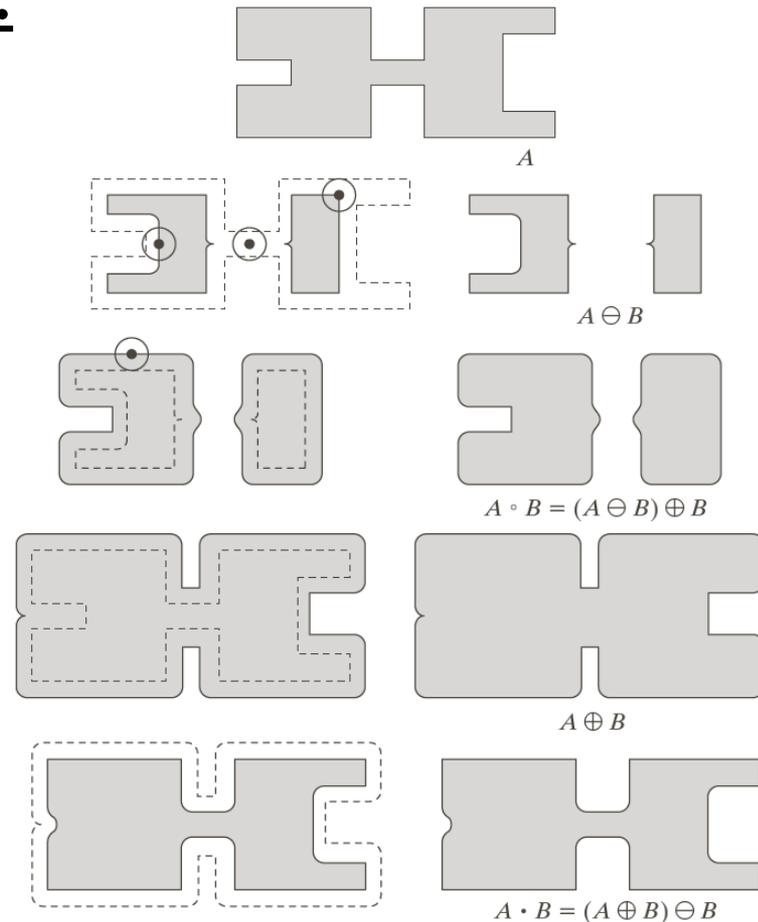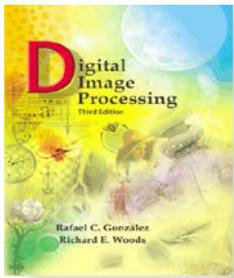
$$A$$

$$A \ominus B$$

$$A \circ B = (A \ominus B) \oplus B$$

$$A \oplus B$$

$$A \cdot B = (A \oplus B) \ominus B$$

Chapter 9
Morphological Image Processing
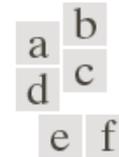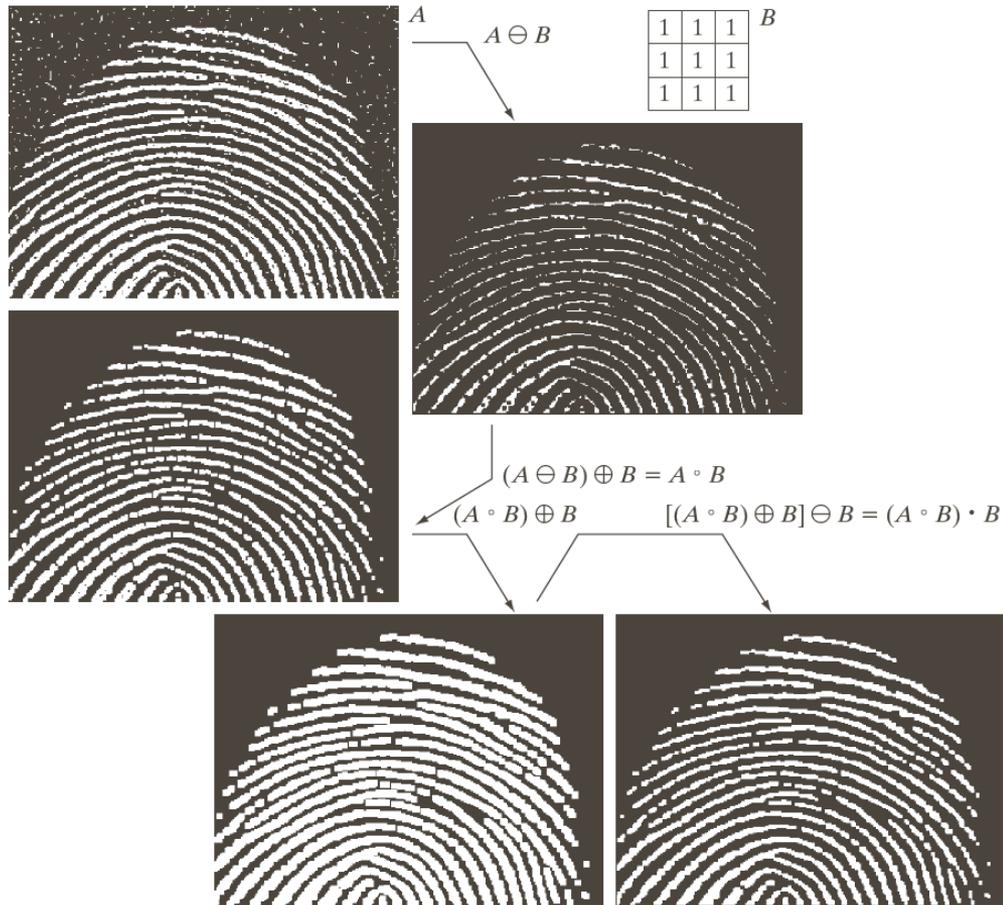
FIGURE 9.11
(a) Noisy image.
(b) Structuring element.
(c) Eroded image.
(d) Opening of *A*.
(e) Dilation of the opening.
(f) Closing of the opening.
(Original image courtesy of the National Institute of Standards and Technology.)

Hit-or-Miss Transformation

$$A \circledast B = (A \ominus D) \cap [A^c \ominus (W - D)]$$
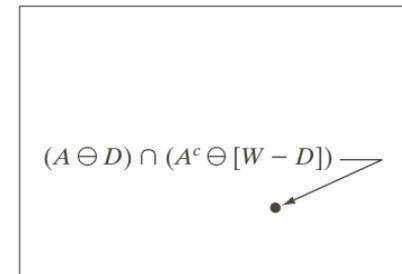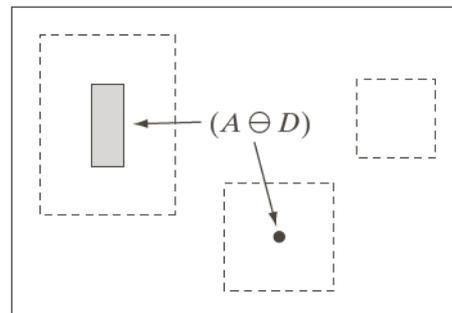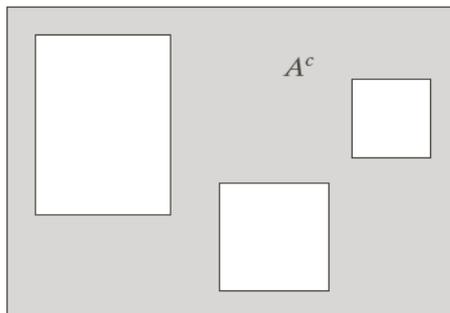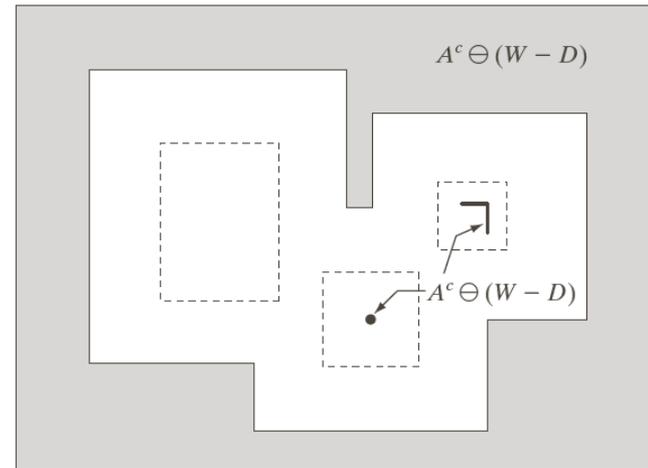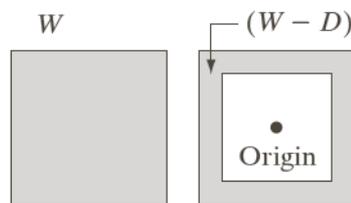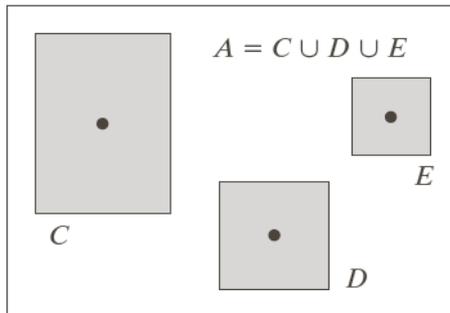
Digital Image Processing
Chapter 11:
Image Description and Representation

Image Representation and Description?

Objective:

   To represent and describe information embedded in an image in other forms that are more suitable than the image itself.

Benefits:

   - Easier to understand
   - Require fewer memory, faster to be processed
   - More "ready to be used"

What kind of information we can use?

   - Boundary, shape
   - Region
   - Texture
   - Relation between regions

Why we focus on a boundary?

The boundary is a good representation of an object shape and also requires a few memory.

Chain codes: represent an object boundary by a connected sequence of straight line segments of specified length and direction.

4-directional
chain code

8-directional
chain code

(Images from Rafael C.
Gonzalez and Richard E.
Wood, Digital Image
Processing, 2nd Edition.
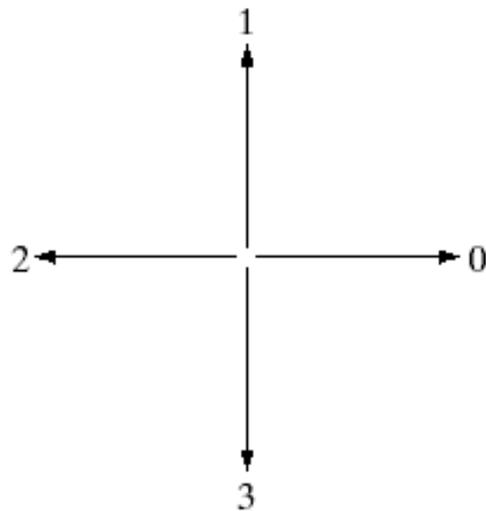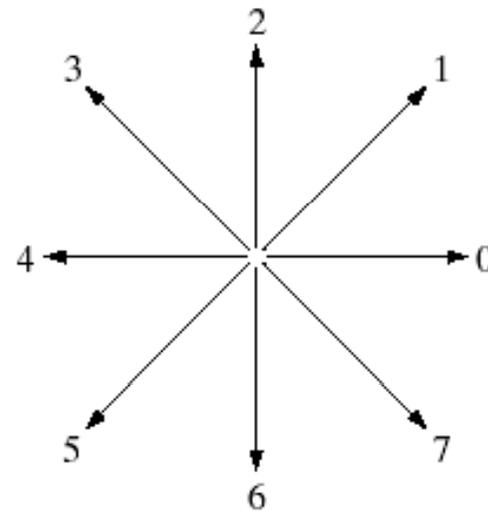
Examples of Chain Codes

Object boundary (resampling)
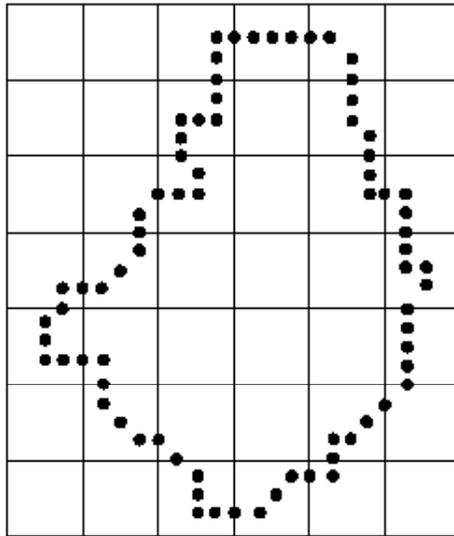
Boundary vertices
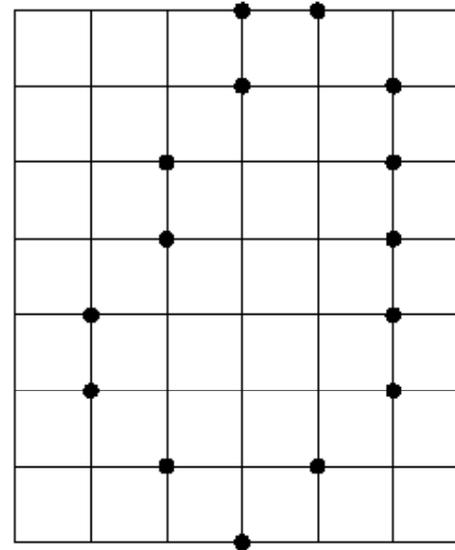
4-directional chain code
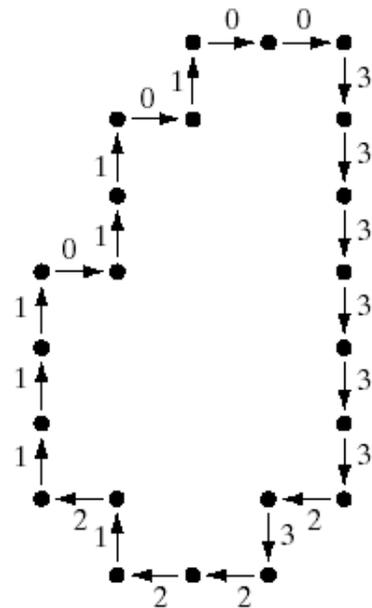
8-directional chain code
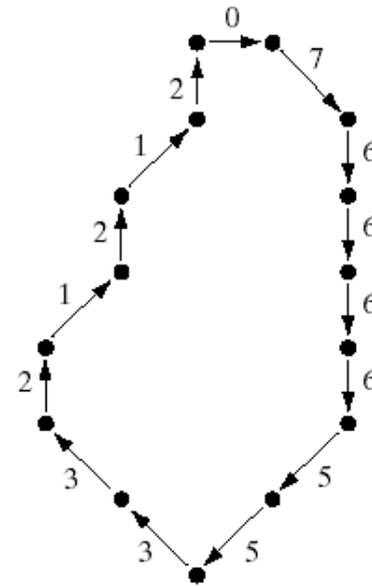
The First Difference of a Chain Codes

Problem of a chain code:
    a chain code sequence depends on a starting point.
Solution: treat a chain code as a circular sequence and redefine the starting point so that the resulting sequence of numbers forms an integer of minimum magnitude.

The first difference of a chain code: counting the number of direction change (in counterclockwise) between 2 adjacent elements of the code.

Example:

Chain code : The first difference
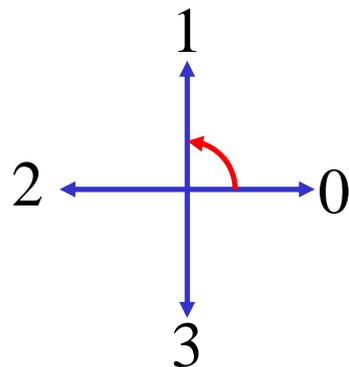
0 → 1	1
0 → 2	2
0 → 3	3
2 → 3	1
2 → 0	2
2 → 1	3

Example:
- a chain code: 10103322
- The first difference = 3133030
- Treating a chain code as a circular sequence, we get the first difference = 33133030

The first difference is rotational invariant.

Polygon Approximation

Represent an object boundary by a polygon
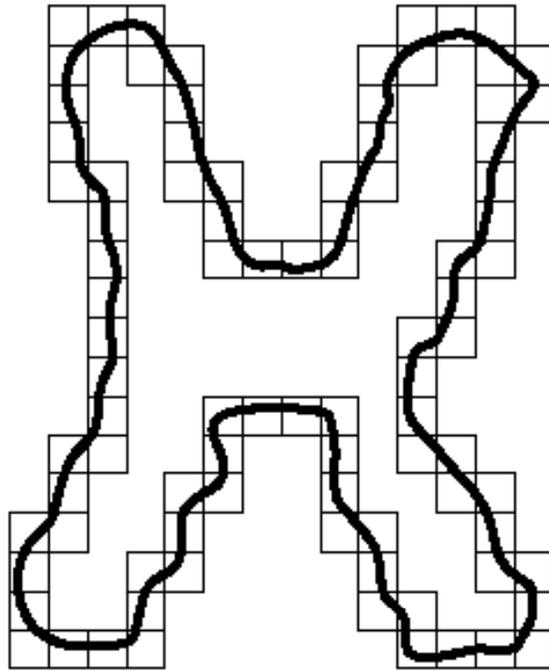
Object boundary
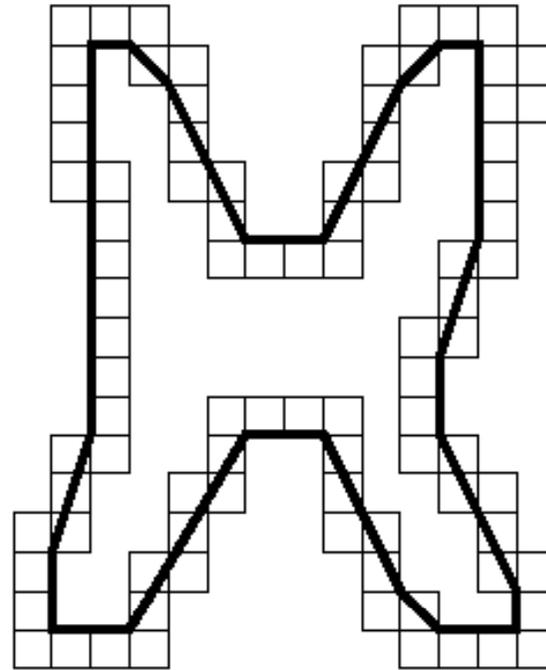
Minimum perimeter polygon

Minimum perimeter polygon consists of line segments that minimize distances between boundary pixels.

Polygon Approximation:Splitting Techniques

0. Object boundary
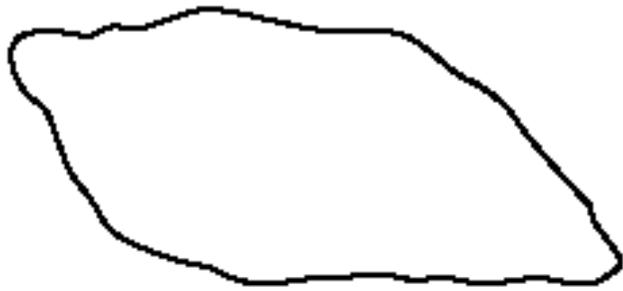
1. Find the line joining two extreme points
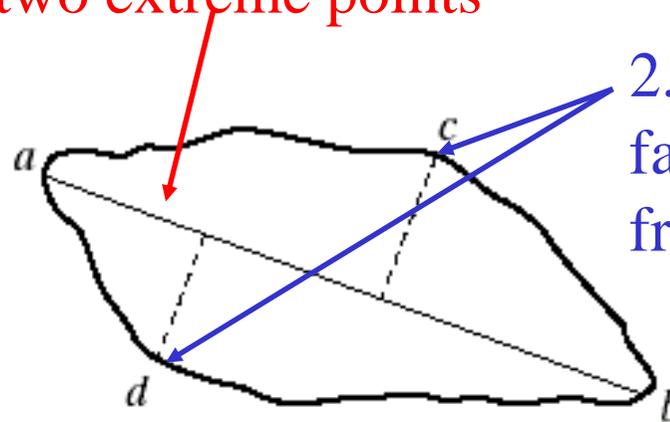
2. Find the farthest points from the line

3. Draw a polygon

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.
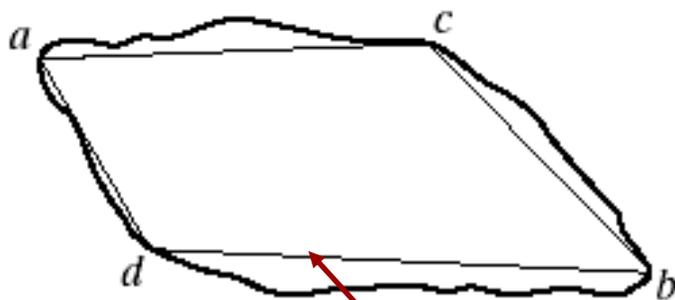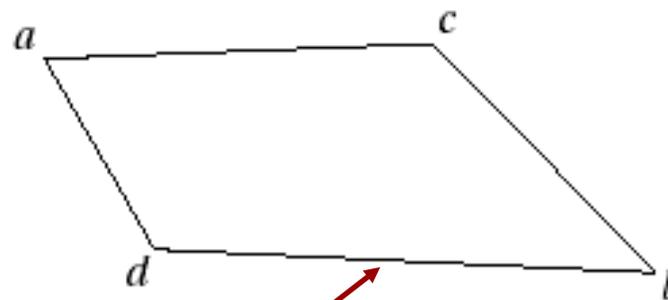
Represent an 2-D object boundary in term of a 1-D function of radial distance with respect to θ.

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.)
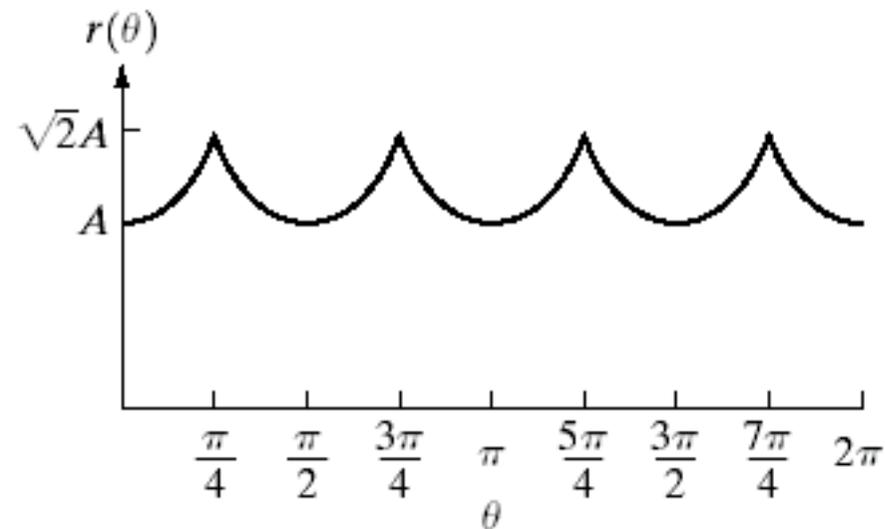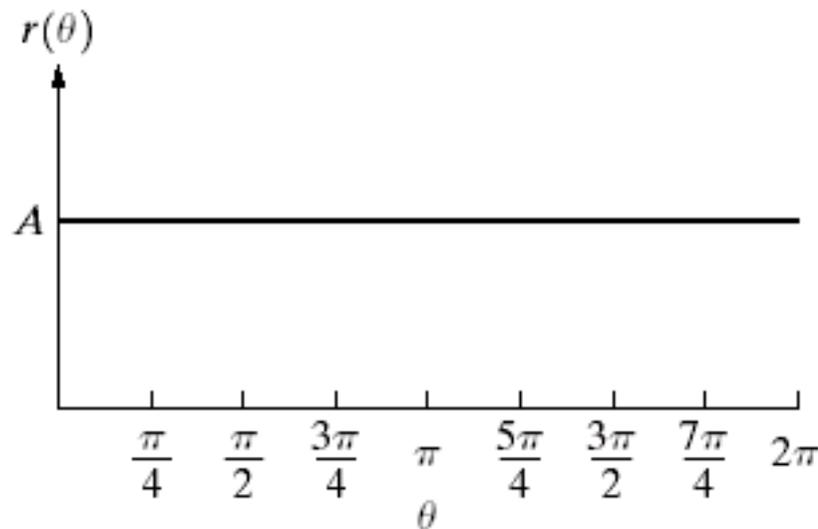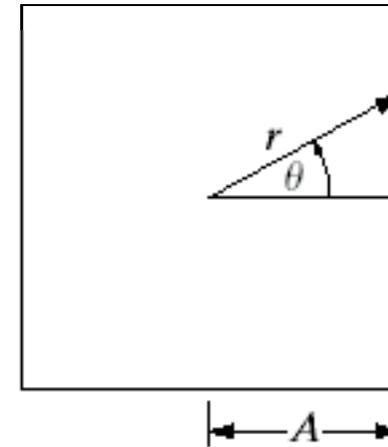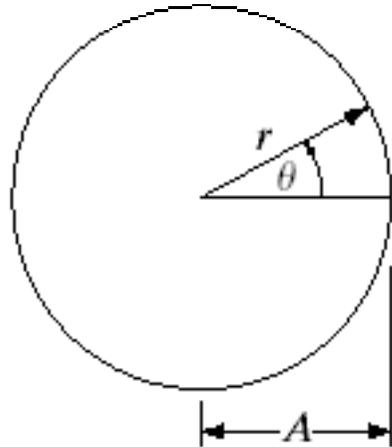
Concept: Partitioning an object boundary by using vertices of a convex hull.

Partitioned boundary

Convex hull (gray color)

Object boundary
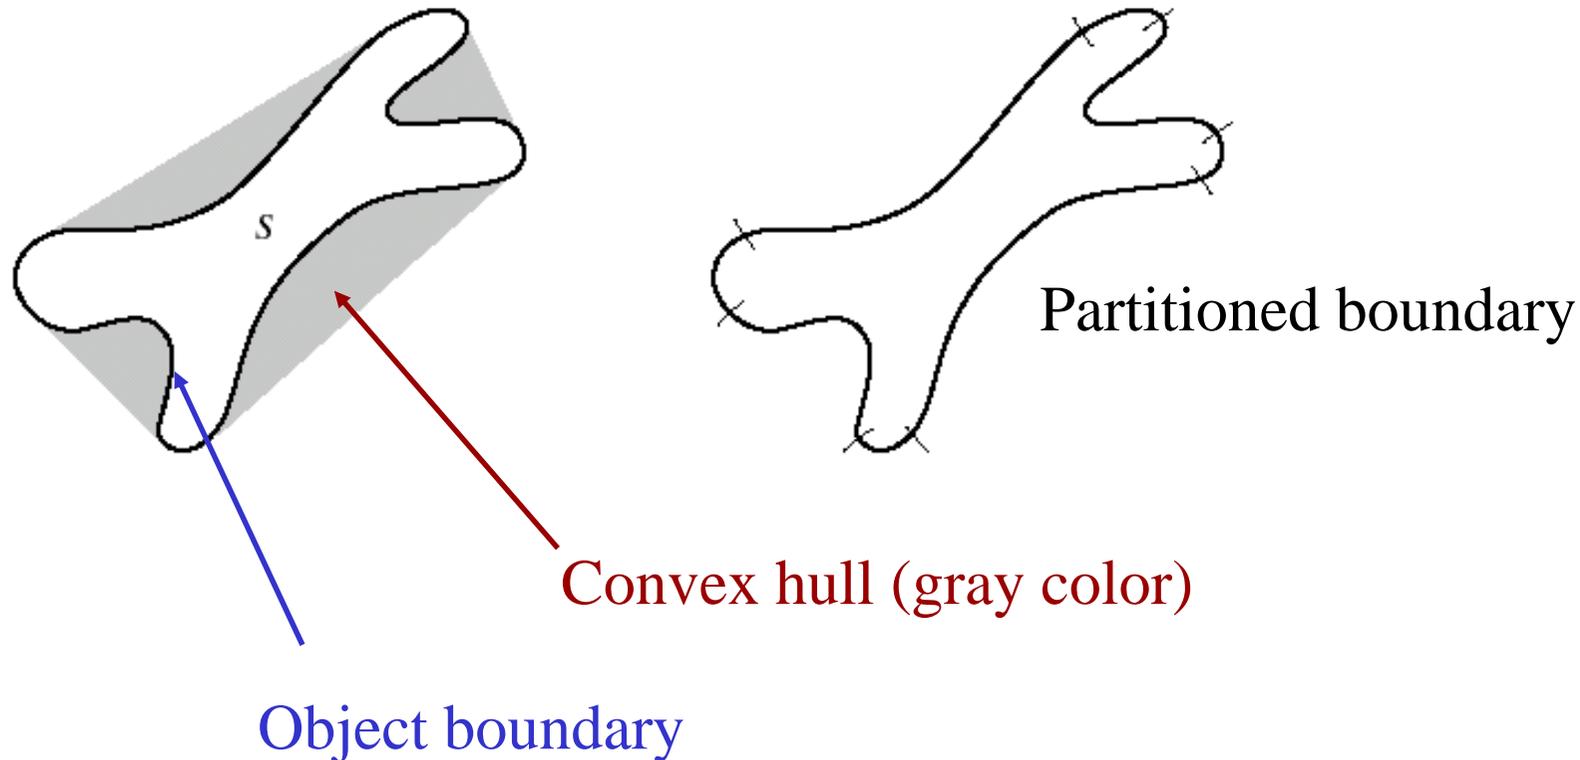
Convex Hull Algorithm

Input   : A set of points on a cornea boundary

Output: A set of points on a boundary of a convex hull of a cornea

1. Sort the points by x-coordinate to get a sequence p_1, p_2, $\bullet\bullet\bullet$,p_n

<u>For the upper side of a convex hull</u>

2. Put the points p_1 and p_2 in a list L_{upper} with p_1 as the first point

3. For i = 3 to n

4.   Do append p_i to L_{upper}

5.   While L_{upper} contains more than 2 points and the last 3

points in L_{upper} do not make a right turn

6.   Do delete the middle point of the last 3 points from L_{upper}
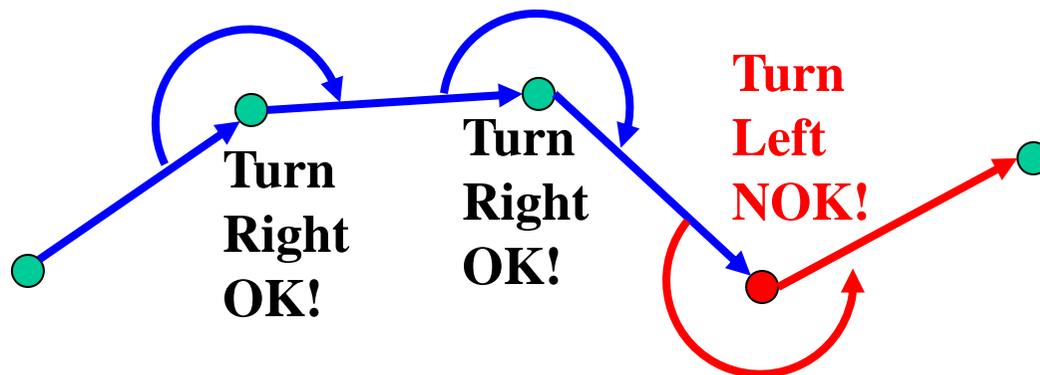
Convex Hull Algorithm (cont.)

For the lower side of a convex hull

7. **Put** the points p_n and p_{n-1} in a list L_{lower} with p_n as the first point

8. **For** i = n-2 down to 1

9. **Do** append p_i to L_{lower}

10. **While** L_{lower} contains more than 2 points and **the last 3 points**

in L_{lower} do not make a right turn

11. **Do delete the middle point** of the last 3 points from L_{lower}

12. **Remove** the first and the last points from L_{lower}

13. **Append** L_{lower} to L_{upper} resulting in the list L

14. **Return** L
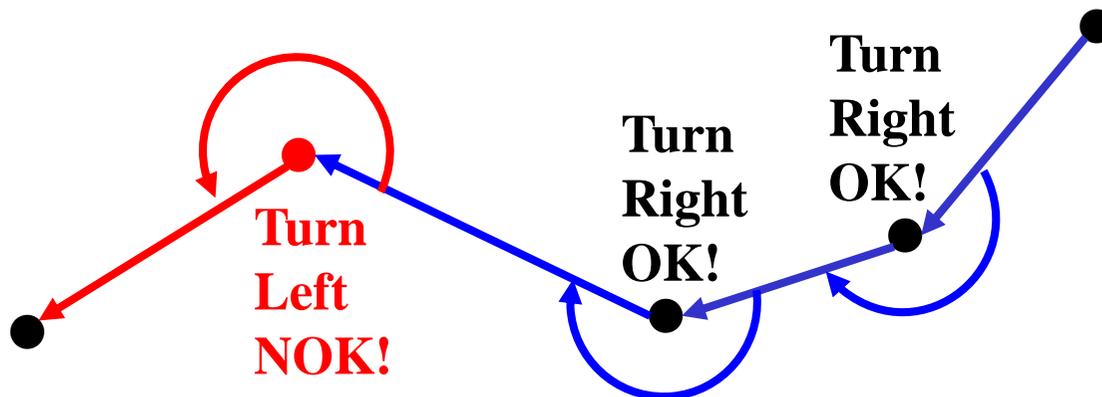
Skeletons

Obtained from thinning or skeletonizing processes
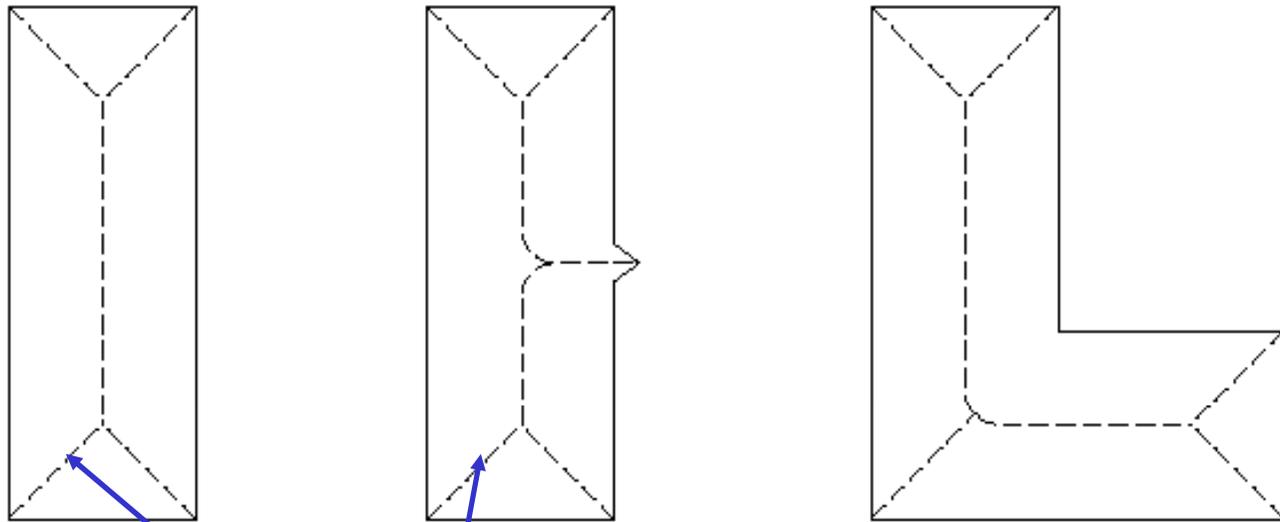
a b c

FIGURE 11.7
Medial axes
(dashed) of three
simple regions.

Medial axes (dash lines)

(Images from Rafael C. Gonzalez and Richard E.
Wood, Digital Image Processing, 2nd Edition.

Thinning Algorithm

Concept: 1. Do not remove end points

                  2. Do not break connectivity

                  3. Do not cause excessive erosion

Apply only to contour pixels: pixels "1" having at least one of its 8 neighbor pixels valued "0"

Notation:

Let

p_9	p_2	p_3
p_8	p_1	p_4
p_7	p_6	p_5

= Neighborhood arrangement for the thinning algorithm
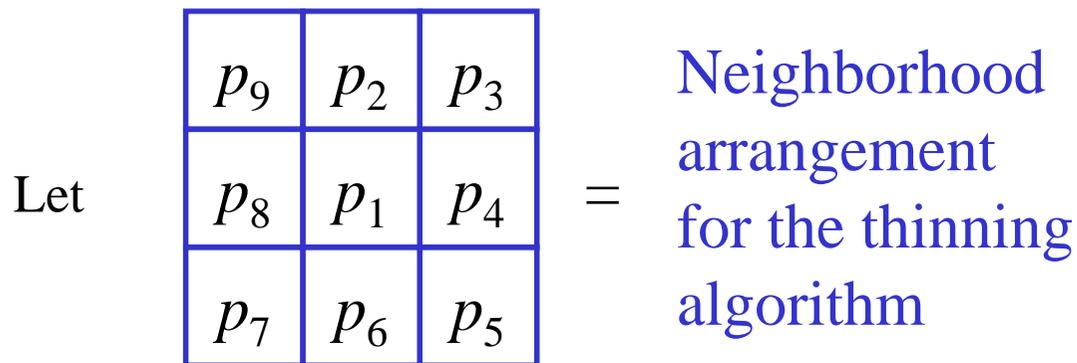
Let
$$N(p_1) = p_2 + p_3 + \ldots + p_8 + p_9$$

$T(p_1)$ = the number of transition 0-1 in the ordered sequence $p_2, p_3, \ldots, p_8, p_9, p_2$.

Example

0	0	1
1	p_1	0
1	0	1

$N(p_1) = 4$
$T(p_1) = 3$

Step 1. Mark pixels for deletion if the following conditions are true.

a) $2 \leq N(p_1) \leq 6$

b) $T(p_1) = 1$ (Apply to all border pixels)

c) $p_2 \cdot p_4 \cdot p_6 = 0$

d) $p_4 \cdot p_6 \cdot p_8 = 0$

p_9	p_2	p_3
p_8	p_1	p_4
p_7	p_6	p_5

Step 2. Delete marked pixels and go to Step 3.

Step 3. Mark pixels for deletion if the following conditions are true.

a) $2 \leq N(p_1) \leq 6$ (Apply to all border pixels)

b) $T(p_1) = 1$

c) $p_2 \cdot p_4 \cdot p_8 = 0$

d) $p_2 \cdot p_6 \cdot p_8 = 0$

Step 4. Delete marked pixels and repeat Step 1 until no change occurs.

Example: Skeletons Obtained from the Thinning Alg.
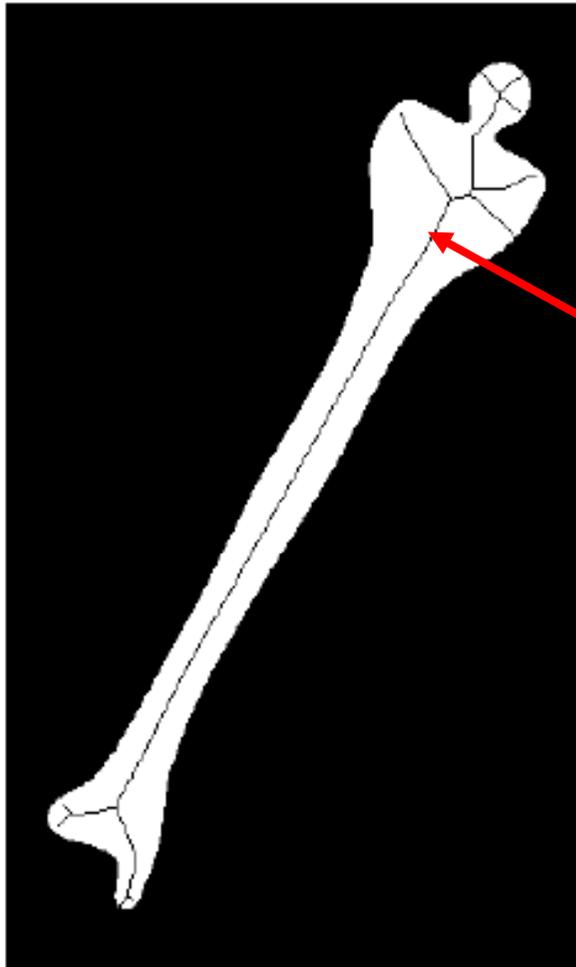
FIGURE 11.10
Human leg bone
and skeleton of
the region shown
superimposed.

Skeleton

Boundary Descriptors

1. Simple boundary descriptors:

we can use

- Length of the boundary
- The size of smallest circle or box that can totally enclosing the object

2. Shape number

3. Fourier descriptor

4. Statistical moments

Shape Number

Shape number of the boundary definition:
    the first difference of smallest magnitude
The order n of the shape number:
    the number of digits in the sequence
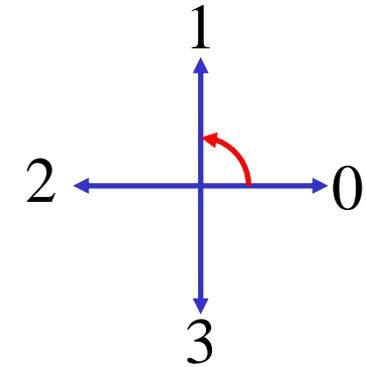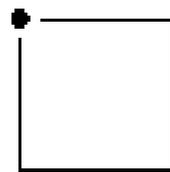
```
         1
         ↑
  2 ←————+————→ 0
         ↓
         3
```

Order 4

Chain code: 0 3 2 1

Difference: 3 3 3 3

Shape no.: 3 3 3 3
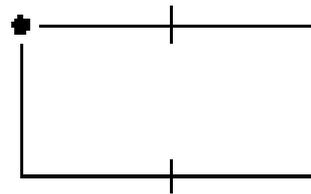
Order 6

Chain code: 0 0 3 2 2 1

Difference: 3 0 3 3 0 3

Shape no.: 0 3 3 0 3 3

Shape numbers of order 4, 6 and 8

Order 4

Chain code: 0 3 2 1

Difference: 3 3 3 3

Shape no.: 3 3 3 3

Order 6

Chain code: 0 0 3 2 2 1

Difference: 3 0 3 3 0 3

Shape no.: 0 3 3 0 3 3

Order 8
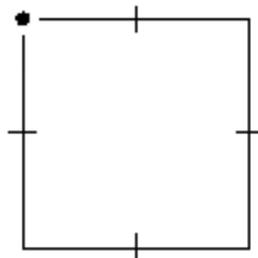
Chain code: 0 0 3 3 2 2 1 1

Difference: 3 0 3 0 3 0 3 0

Shape no.: 0 3 0 3 0 3 0 3
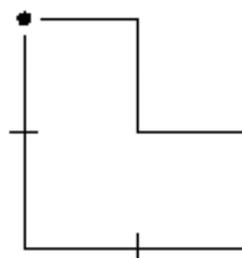
Chain code: 0 3 0 3 2 2 1 1

Difference: 3 3 1 3 3 0 3 0

Shape no.: 0 3 0 3 3 1 3 3

Chain code: 0 0 0 3 2 2 2 1

Difference: 3 0 0 3 3 0 0 3

Shape no.: 0 0 3 3 0 0 3 3

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Example: Shape Number

1. Original boundary
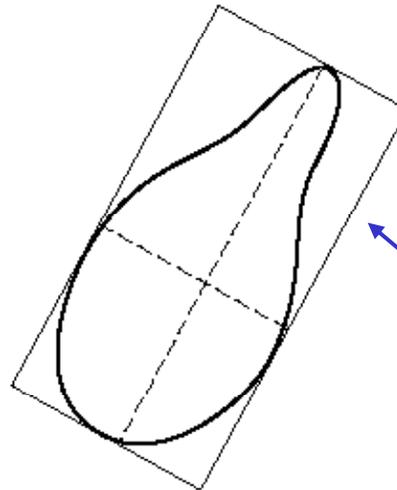
2. Find the smallest rectangle that fits the shape
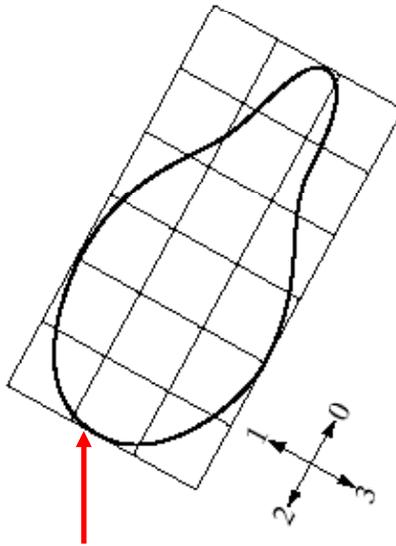
3. Create grid
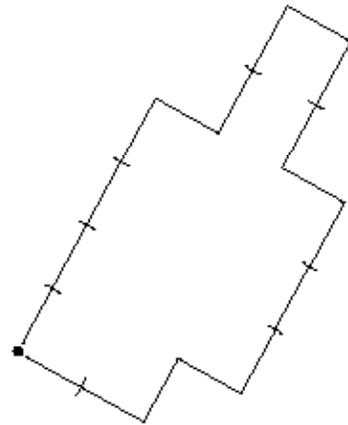
4. Find the nearest Grid.

Chain code:
0 0 0 0 3 0 0 3 2 2 3 2 2 2 1 2 1 1

First difference:
3 0 0 0 3 1 0 3 3 0 1 3 0 0 3 1 3 0

Shape No.
0 0 0 3 1 0 3 3 0 1 3 0 0 3 1 3 0 3

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Fourier Descriptor

Fourier descriptor: view a coordinate (x,y) as a complex number (x = real part and y = imaginary part) then apply the Fourier transform to a sequence of boundary points.

Let s(k) be a coordinate of a boundary point k :

$$s(k) = x(k) + jy(k)$$

Fourier descriptor :

$$a(u) = \frac{1}{K}\sum_{k=0}^{K-1} s(k)e^{-2\pi uk/K}$$

Reconstruction formula

$$s(k) = \frac{1}{K}\sum_{k=0}^{K-1} a(u)e^{2\pi uk/K}$$

Boundary points
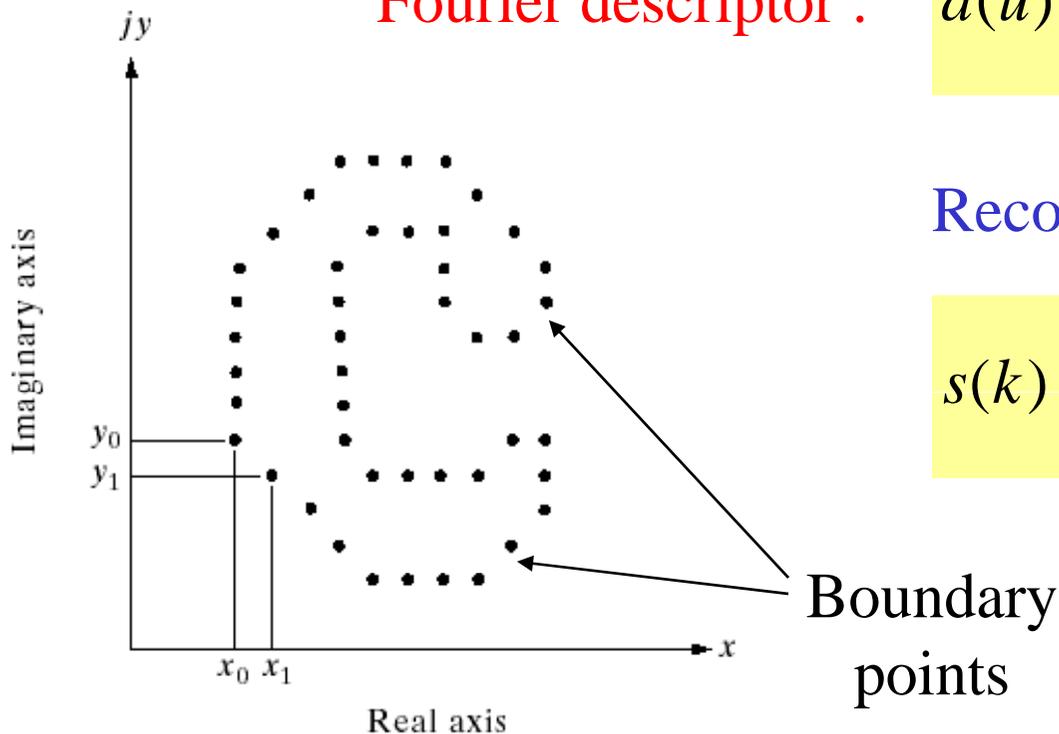
Example: Fourier Descriptor

Examples of reconstruction from Fourier descriptors

Original ($K = 64$) $P = 2$ $P = 4$ $P = 8$

$P = 16$ $P = 24$ $P = 32$ $P = 40$

$P = 48$ $P = 56$ $P = 61$ $P = 62$
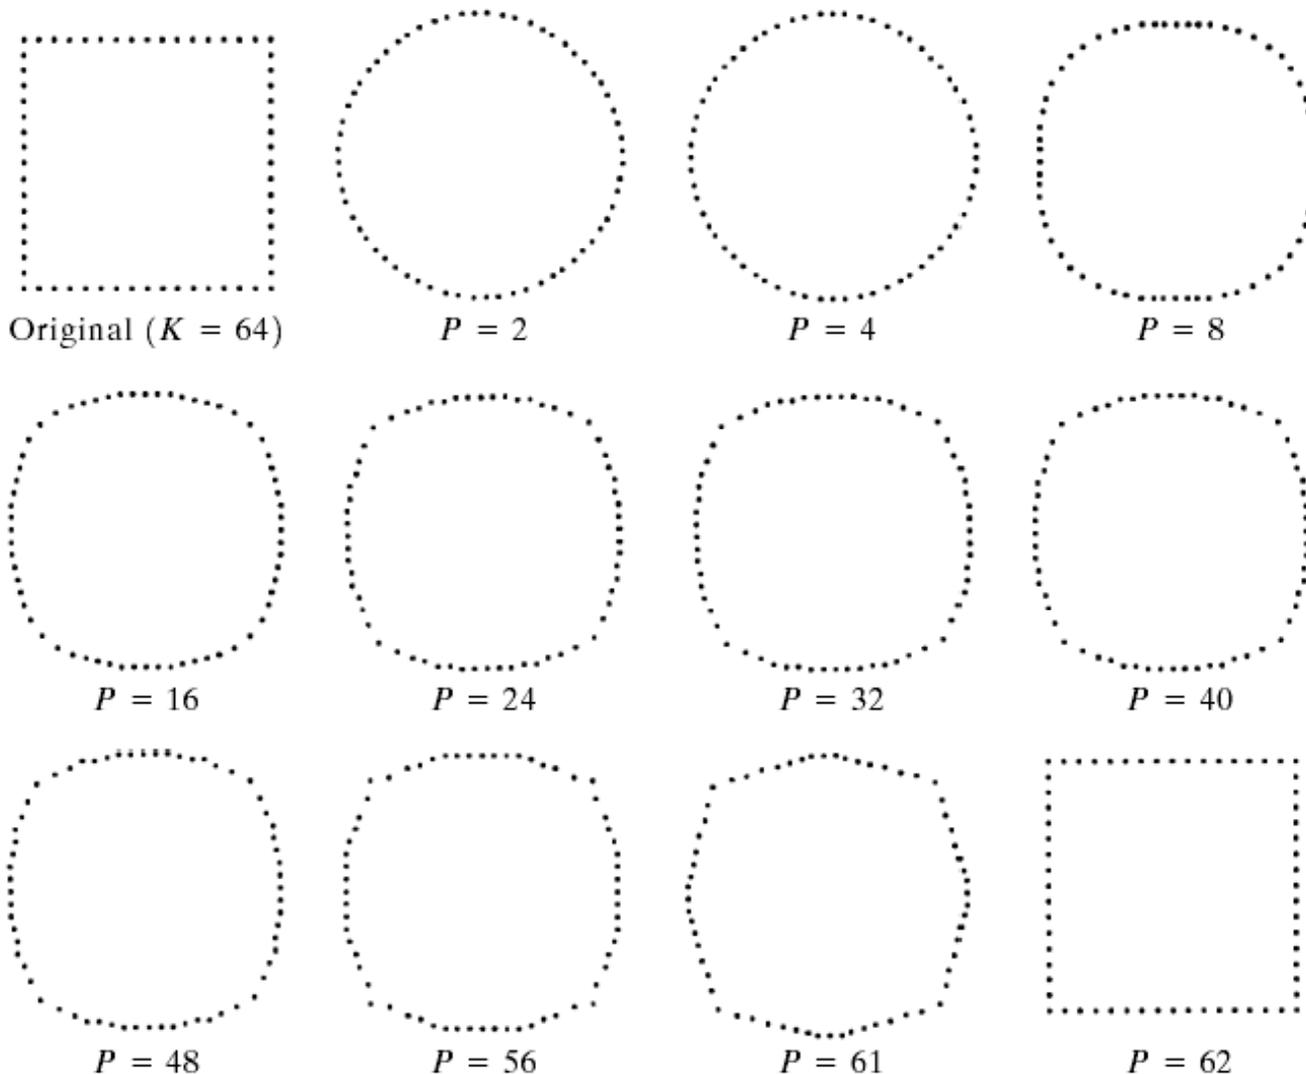
$$\hat{s}(k) = \frac{1}{K} \sum_{k=0}^{P-1} a(u) e^{2\pi u k / K}$$

P is the number of Fourier coefficients used to reconstruct the boundary

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Fourier Descriptor Properties

Some properties of Fourier descriptors

Transformation	Boundary	Fourier Descriptor
Identity	$s(k)$	$a(u)$
Rotation	$s_r(k) = s(k)e^{j\theta}$	$a_r(u) = a(u)e^{j\theta}$
Translation	$s_t(k) = s(k) + \Delta_{xy}$	$a_t(u) = a(u) + \Delta_{xy}\delta(u)$
Scaling	$s_s(k) = \alpha s(k)$	$a_s(u) = \alpha a(u)$
Starting point	$s_p(k) = s(k - k_0)$	$a_p(u) = a(u)e^{-j2\pi k_0 u/K}$

Statistical Moments

Definition: the n^{th} moment

$$\mu_n(r) = \sum_{i=0}^{K-1} (r_i - m)^n g(r_i)$$

where

$$m = \sum_{i=0}^{K-1} r_i g(r_i)$$

Example of moment:
The first moment = mean
The second moment = variance

Boundary segment

$g(r)$
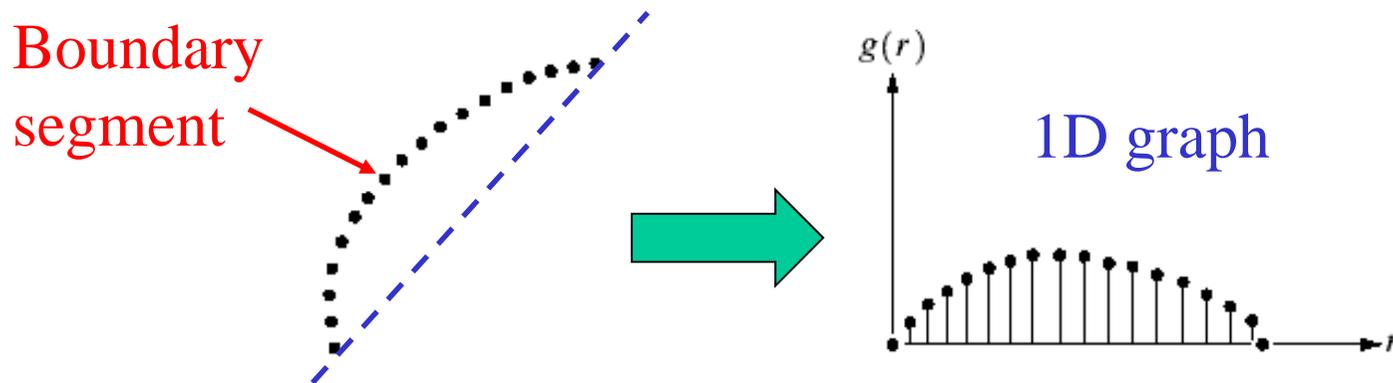
1D graph

1. Convert a boundary segment into 1D graph
2. View a 1D graph as a PDF function
3. Compute the n^{th} order moment of the graph

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.)

Regional Descriptors

Purpose: to describe regions or "areas"

1. Some simple regional descriptors
- area of the region
- length of the boundary (perimeter) of the region
- Compactness

$$C = \frac{A(R)}{P^2(R)}$$

where $A(R)$ and $P(R)$ = area and perimeter of region R

Example: a circle is the most compact shape with $C = 1/4\pi$

2. Topological Descriptors
3. Texture
4. Moments of 2D Functions

Example: Regional Descriptors
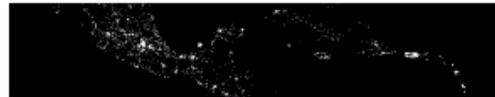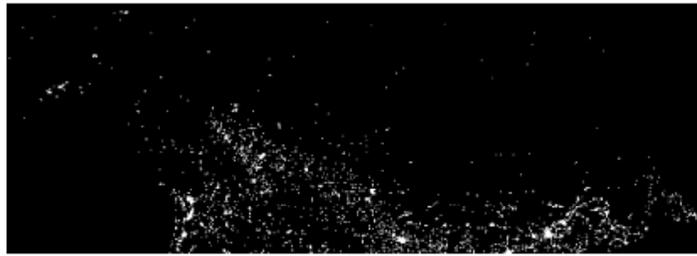
White pixels represent "light of the cities"

Region no.	% of white pixels compared to the total white pixels
1	20.4%
2	64.0%
3	4.9%
4	10.7%

Infrared image of America at night

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Topological Descriptors

Use to describe holes and connected components of the region
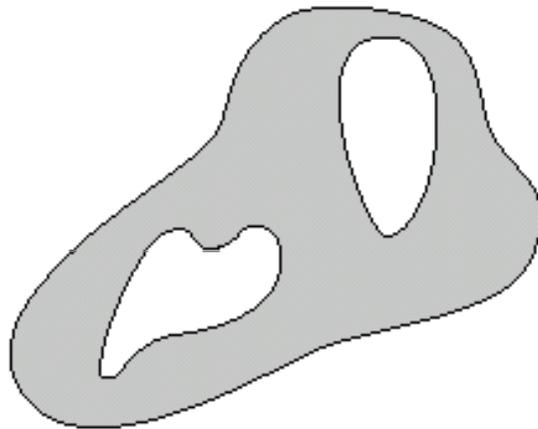
FIGURE 11.17 A region with two holes.

Euler number (E):

$$E = C - H$$

C = the number of connected components

H = the number of holes
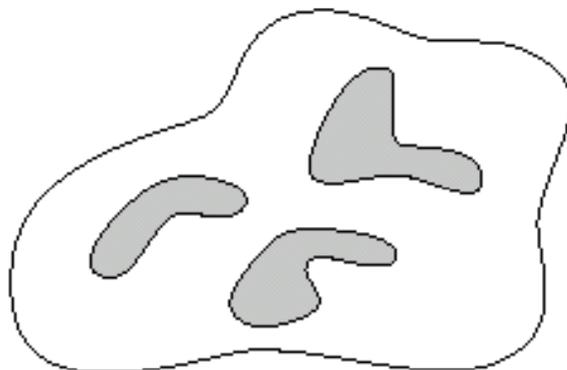
FIGURE 11.18 A region with three connected components.

Topological Descriptors (cont.)

E = -1

E = 0

E = -2
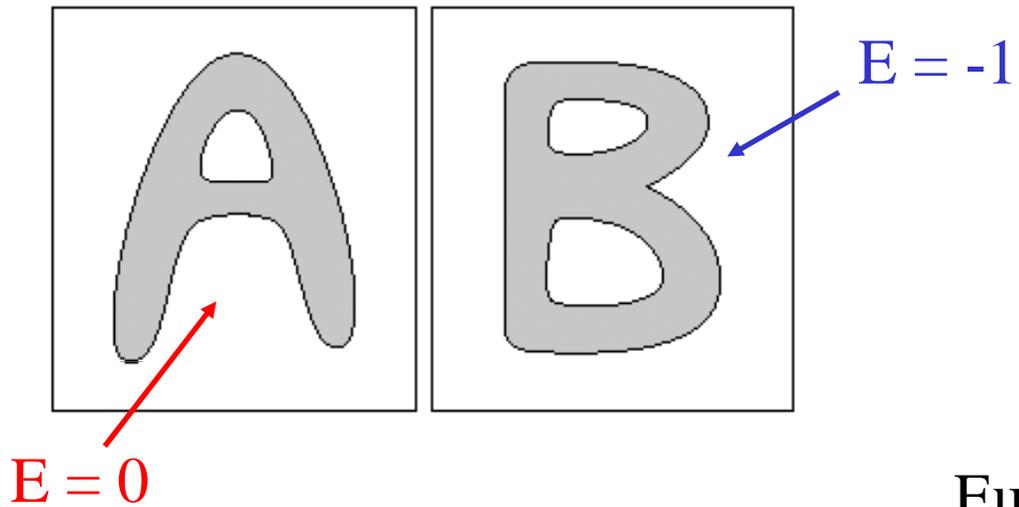
Euler Formula

$$V - Q + F = C - H = E$$
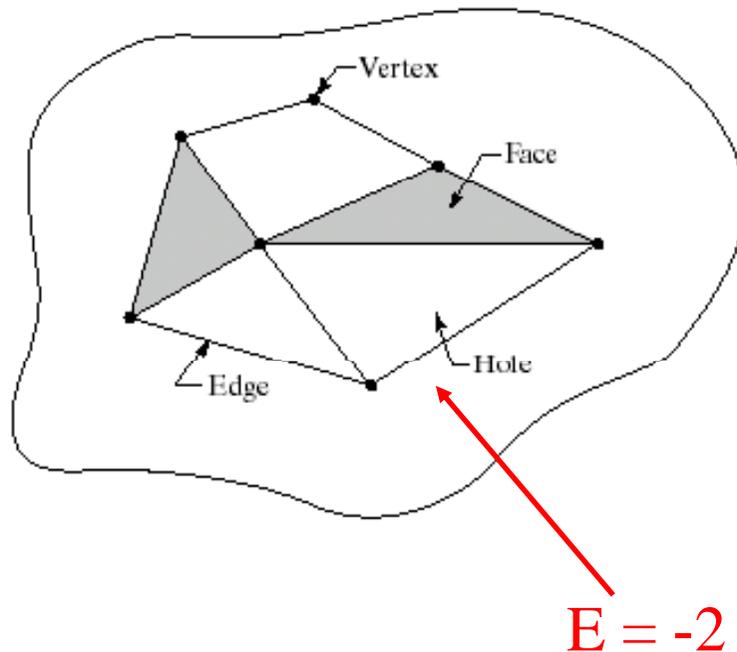
V = the number of vertices
Q = the number of edges
F = the number of faces

Vertex

Face

Edge

Hole

Example: Topological Descriptors

Original image:
Infrared image
Of Washington
D.C. area
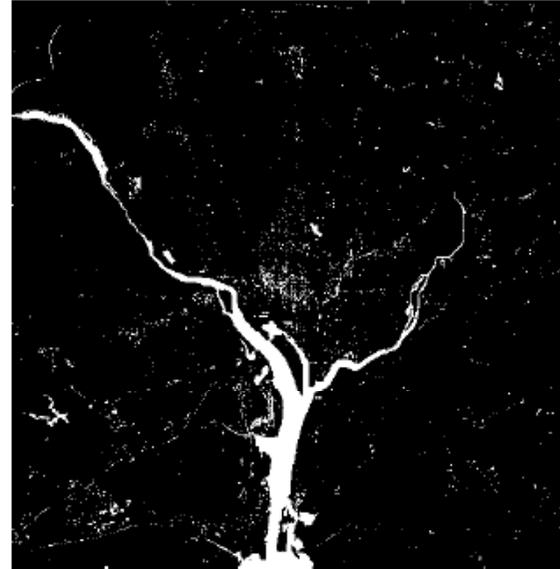
After intensity
Thresholding
(1591 connected
components
with 39 holes)
Euler no. = 1552

The largest
connected
area
(8479 Pixels)
(Hudson river)

After thinning

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.
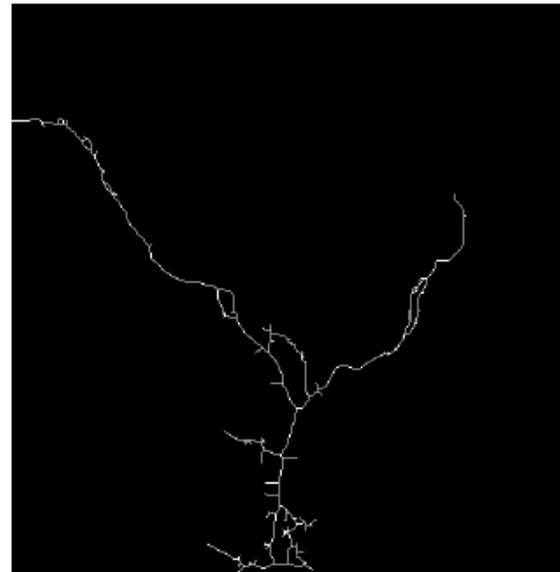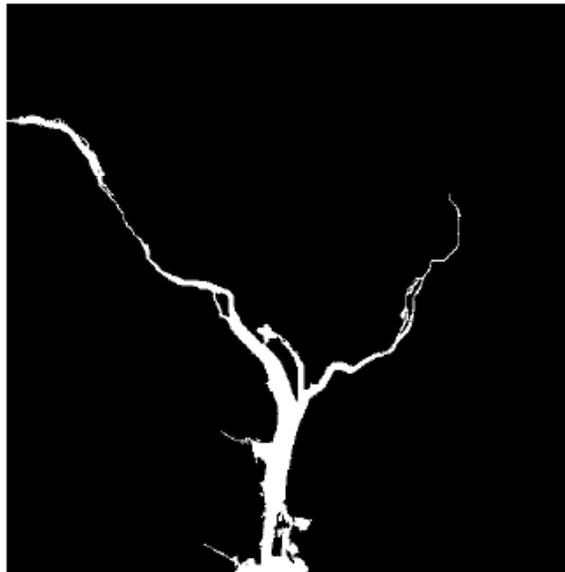
Texture Descriptors

Purpose: to describe "texture" of the region.

Examples: optical microscope images:
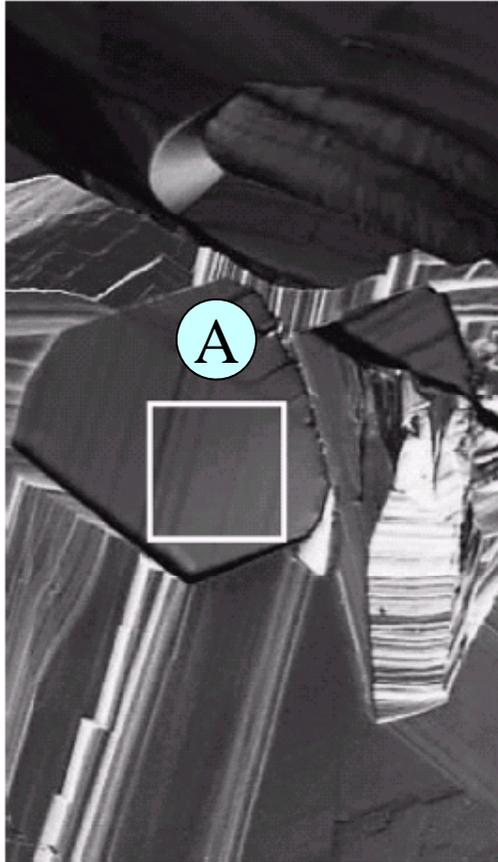
Superconductor
(smooth texture)
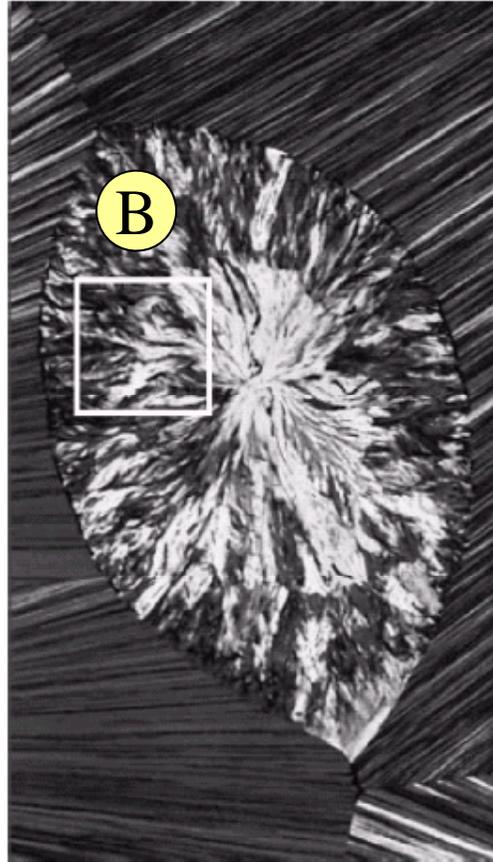
Cholesterol
(coarse texture)
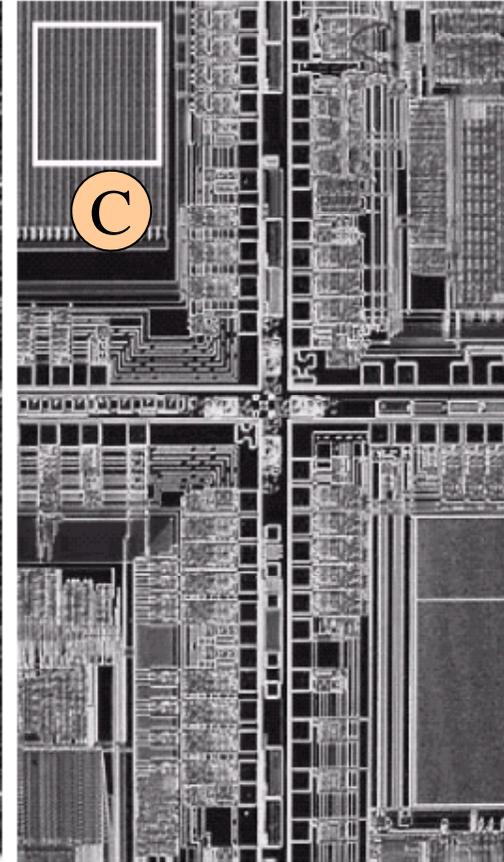
Microprocessor
(regular texture)

Statistical Approaches for Texture Descriptors

We can use statistical moments computed from an image histogram:

$$\mu_n(z) = \sum_{i=0}^{K-1}(z_i - m)^n \, p(z_i)$$

$z =$ intensity

$p(z) -$ PDF or histogram of z

where

$$m = \sum_{i=0}^{K-1} z_i \, p(z_i)$$

Example: The 2nd moment = variance $\rightarrow$ measure "smoothness"

The 3rd moment $\rightarrow$ measure "skewness"

The 4th moment $\rightarrow$ measure "uniformity" (flatness)

Texture	Mean	Standard deviation	R (normalized)	Third moment	Uniformity	Entropy
(A) Smooth	82.64	11.79	0.002	−0.105	0.026	5.434
(B) Coarse	143.56	74.63	0.079	−0.151	0.005	7.783
(C) Regular	99.72	33.73	0.017	0.750	0.013	6.674

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Fourier Approach for Texture Descriptor

Concept: convert 2D spectrum into 1D graphs

Original image → FFT2D +FFTSHIFT → Fourier coefficient image → Divide into areas by angles

Fourier coefficient image ↓ Divide into areas by radius

Divide into areas by angles ↓ Sum all pixels in each area

$$S(\theta) = \sum_{r=1}^{R_0} S_r(\theta)$$

Divide into areas by radius ↓ Sum all pixels in each area →

$$S(r) = \sum_{\theta=0}^{\pi} S_\theta(r)$$
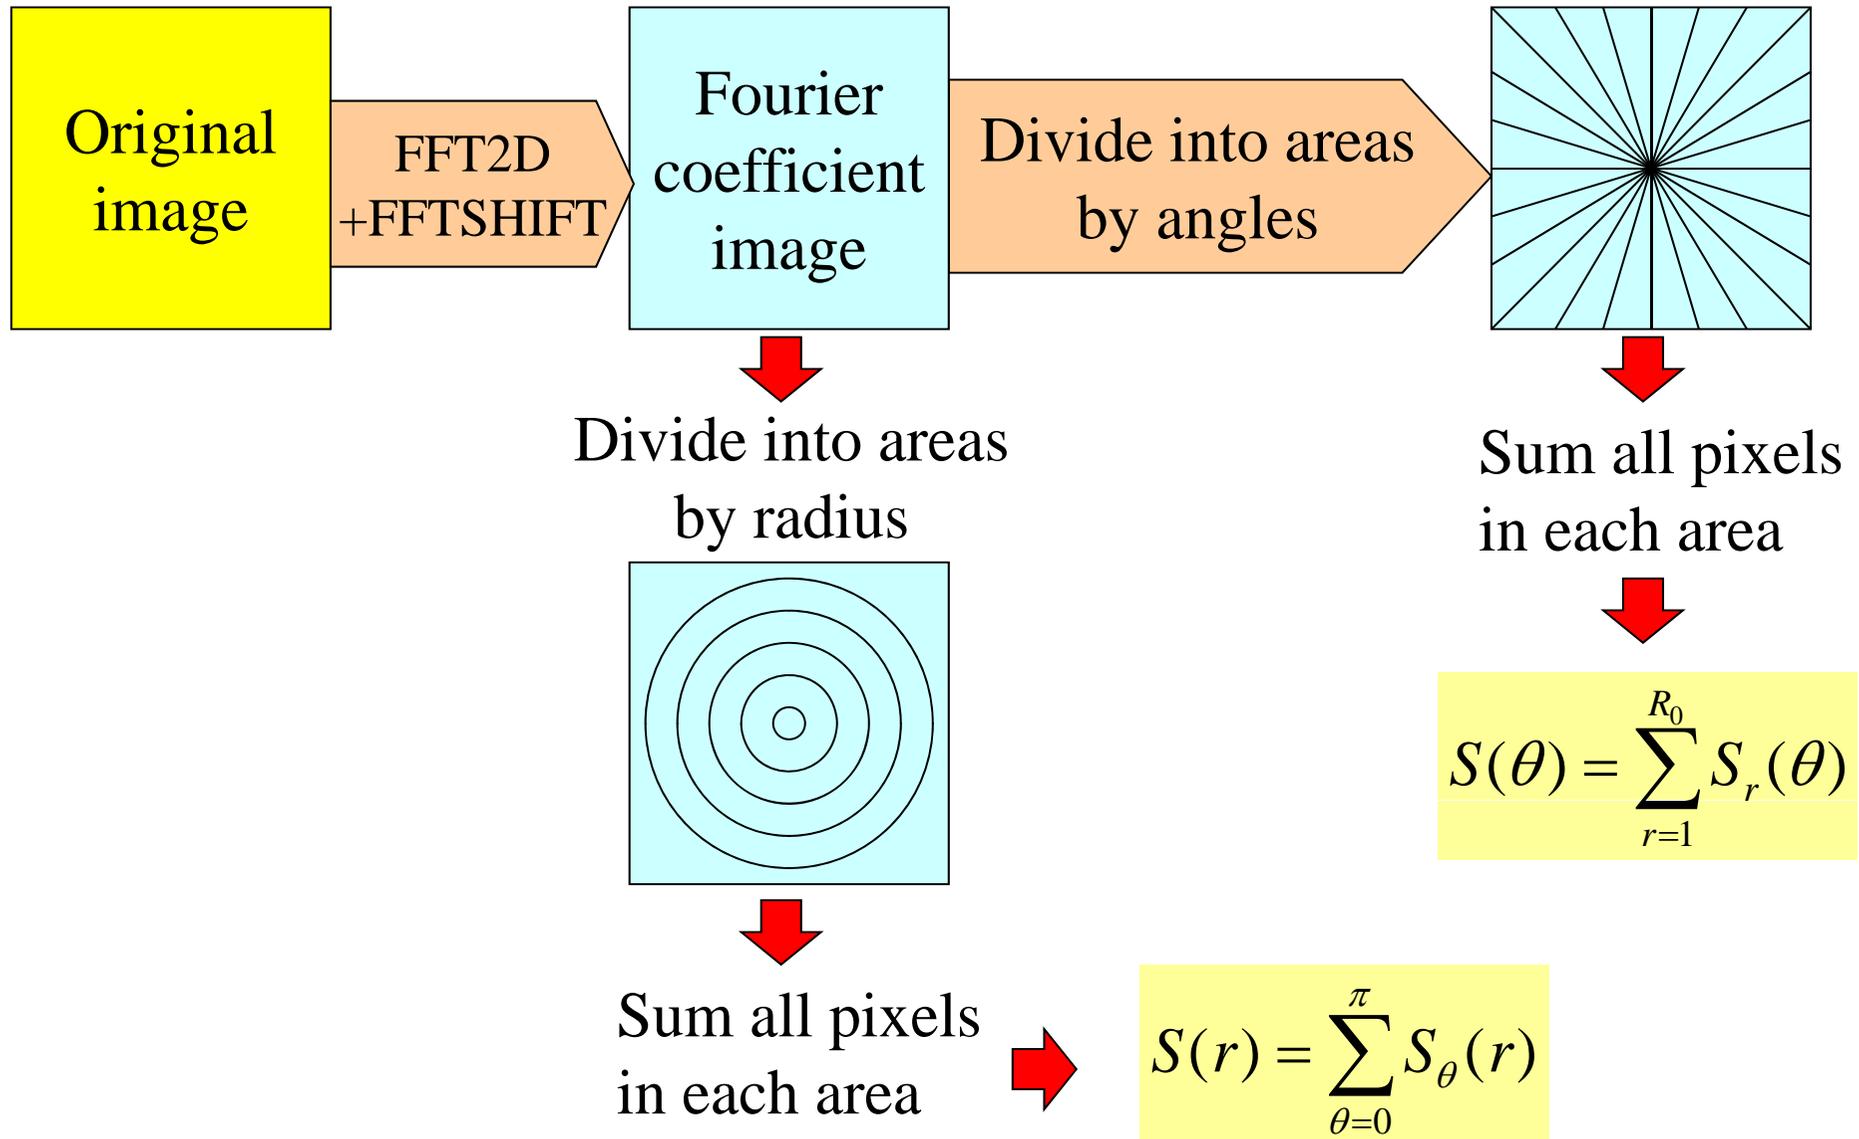
Fourier Approach for Texture Descriptor

Original image
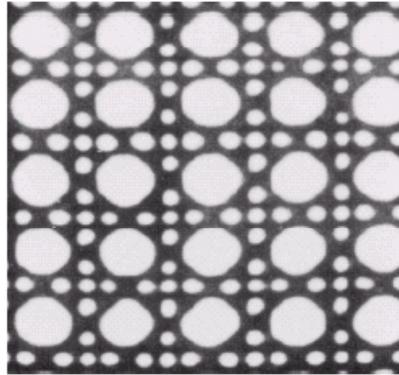
2D Spectrum (Fourier Tr.)
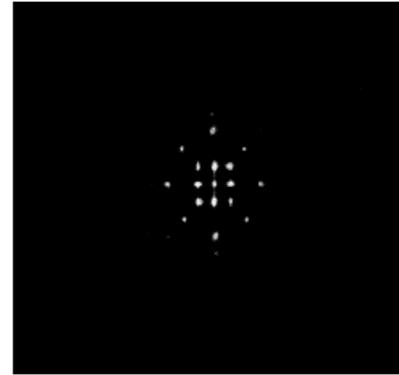
$S(r)$
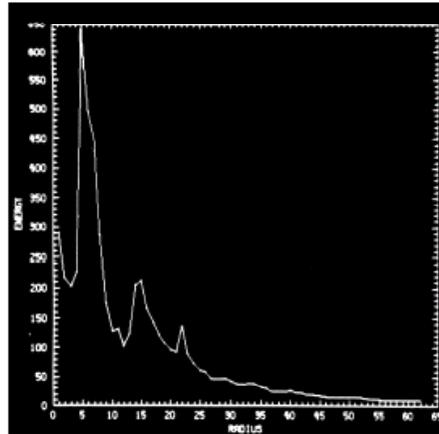
$S(\theta)$
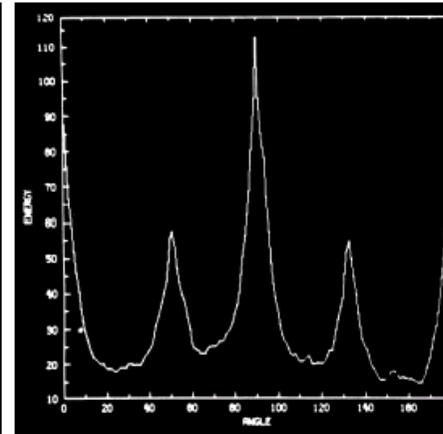
Another image
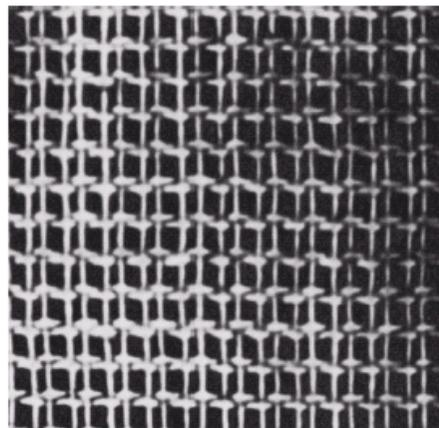
Another $S(\theta)$
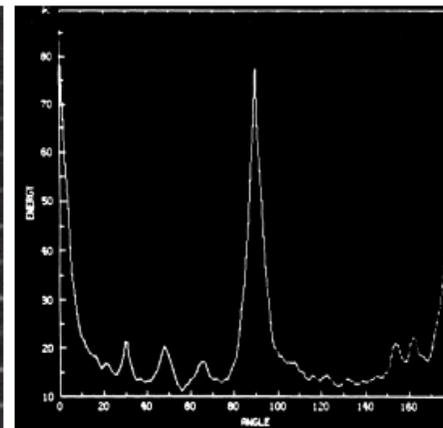
(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Moments of Two-D Functions

The moment of order $p + q$

$$m_{pq} = \sum_x \sum_y x^p y^q f(x, y)$$

$$\bar{x} = \frac{m_{10}}{m_{00}}$$

$$\bar{y} = \frac{m_{01}}{m_{00}}$$

The central moments of order $p + q$

$$\mu_{pq} = \sum_x \sum_y (x - \bar{x})^p (y - \bar{y})^q f(x, y)$$

$$\mu_{00} = m_{00} \qquad \mu_{01} = \mu_{10} = 0$$

$$\mu_{11} = m_{11} - \bar{x}m_{01} = m_{11} - \bar{y}m_{10}$$

$$\mu_{20} = m_{20} - \bar{x}m_{10} \qquad \mu_{02} = m_{02} - \bar{y}m_{01}$$

$$\mu_{21} = m_{21} - 2\bar{x}m_{11} - \bar{y}m_{20} + 2\bar{x}^2 m_{01} \qquad \mu_{30} = m_{30} - 3\bar{x}m_{20} + 2\bar{x}^2 m_{10}$$

$$\mu_{12} = m_{12} - 2\bar{y}m_{11} - \bar{x}m_{02} + 2\bar{y}^2 m_{10} \qquad \mu_{03} = m_{03} - 3\bar{y}m_{02} + 2\bar{y}^2 m_{01}$$

Invariant Moments of Two-D Functions

The normalized central moments of order $p + q$

$$\eta_{pq} = \frac{\mu_{pq}}{\mu_{00}^{\gamma}} \qquad \text{where} \qquad \gamma = \frac{p+q}{2} + 1$$

Invariant moments: independent of rotation, translation, scaling, and reflection

$$\phi_1 = \eta_{20} + \eta_{02} \qquad \phi_2 = (\eta_{20} - \eta_{02})^2 + 4\eta_{11}^2$$

$$\phi_3 = (\eta_{30} - 3\eta_{12})^2 + (3\eta_{21} - \eta_{03})^2 \qquad \phi_4 = (\eta_{30} + \eta_{12})^2 + (\eta_{21} + \eta_{03})^2$$

$$\phi_5 = (\eta_{30} - 3\eta_{12})(\eta_{30} + \eta_{12})\left[(\eta_{30} + \eta_{12})^2 - 3(\eta_{21} + \eta_{03})^2\right] + (3\eta_{21} - \eta_{03})(\eta_{21} + \eta_{03})\left[3(\eta_{30} + \eta_{12})^2 - (\eta_{21} + \eta_{03})^2\right]$$

$$\phi_6 = (\eta_{20} - \eta_{02})\left[(\eta_{30} + \eta_{12})^2 - (\eta_{21} + \eta_{03})^2\right] + 4\eta_{11}(\eta_{30} + \eta_{12})(\eta_{21} + \eta_{03})$$

Example: Invariant Moments of Two-D Functions

1. Original image
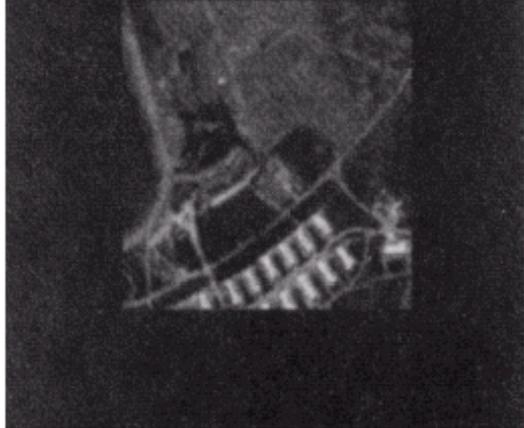
2. Half size

3. Mirrored
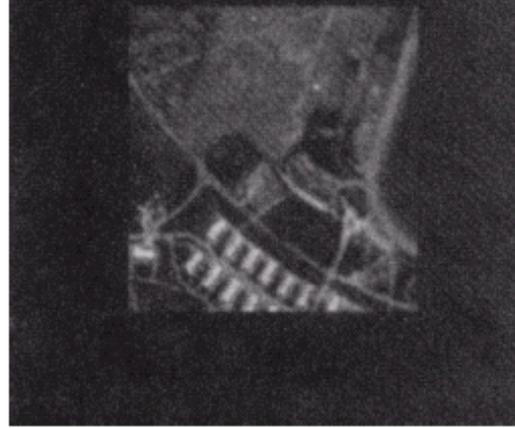
4. Rotated 2 degree

5. Rotated 45 degree

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Example: Invariant Moments of Two-D Functions

Invariant moments of images in the previous slide

Invariant (Log)	Original	Half Size	Mirrored	Rotated 2°	Rotated 45°
ϕ_1	6.249	6.226	6.919	6.253	6.318
ϕ_2	17.180	16.954	19.955	17.270	16.803
ϕ_3	22.655	23.531	26.689	22.836	19.724
ϕ_4	22.919	24.236	26.901	23.130	20.437
ϕ_5	45.749	48.349	53.724	46.136	40.525
ϕ_6	31.830	32.916	37.134	32.068	29.315
ϕ_7	45.589	48.343	53.590	46.017	40.470

Invariant moments are independent of rotation, translation, scaling, and reflection

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Principal Components for Description

Purpose: to reduce dimensionality of a vector image while maintaining information as much as possible.

Let $\quad \mathbf{x} = [x_1 \quad x_2 \quad ... \quad x_n]^T$

Mean: $\quad \mathbf{m}_x = E\{\mathbf{x}\} = \dfrac{1}{K}\displaystyle\sum_{k=1}^{K}\mathbf{x}_k$

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix}$$

Spectral band 6

Spectral band 5

Spectral band 4

Spectral band 3

Spectral band 2

Spectral band 1
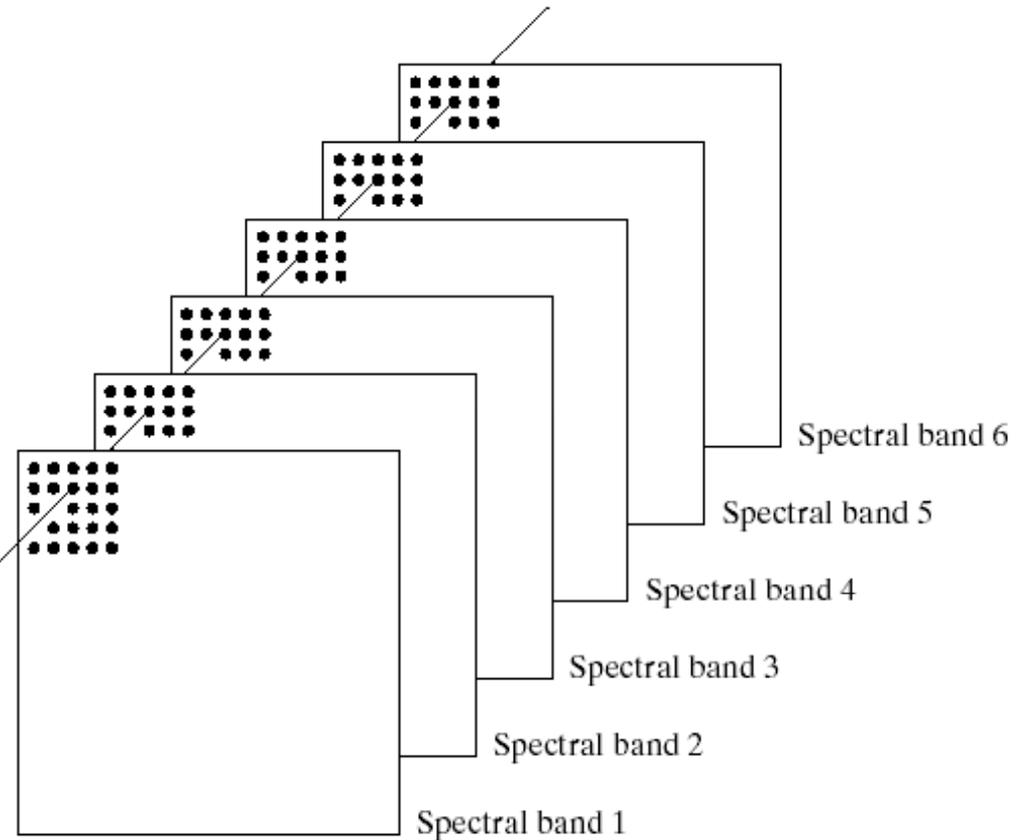
Covariance matrix $\quad \mathbf{C}_x = E\{(\mathbf{x} - \mathbf{m}_x)(\mathbf{x} - \mathbf{m}_x)^T\} = \dfrac{1}{K}\displaystyle\sum_{k=1}^{K}\mathbf{x}_k\mathbf{x}_k^T - \mathbf{m}_x\mathbf{m}_x^T$

Hotelling transformation

Let $\mathbf{y} = \mathbf{A}(\mathbf{x} - \mathbf{m}_x)$

Where A is created from eigenvectors of C_x as follows
Row 1 contain the 1st eigenvector with the largest eigenvalue.
Row 2 contain the 2nd eigenvector with the 2nd largest eigenvalue.
....

Then we get

$$\mathbf{m}_y = E\{\mathbf{y}\} = 0$$

and

$$\mathbf{C}_y = \mathbf{A}\mathbf{C}_x\mathbf{A}^T$$

$$\mathbf{C}_y = \begin{bmatrix} \lambda_1 & 0 & ... & 0 \\ 0 & \lambda_1 & ... & 0 \\ ... & ... & ... & ... \\ 0 & ... & ... & \lambda_1 \end{bmatrix}$$

Then elements of $\mathbf{y} = \mathbf{A}(\mathbf{x} - \mathbf{m}_x)$ are uncorrelated. The component of y with the largest λ is called the principal component.

Eigenvector and Eigenvalue

Eigenvector and eigenvalue of Matrix **C** are defined as

Let **C** be a matrix of size NxN and **e** be a vector of size Nx1. If

$$\mathbf{Ce} = \lambda\mathbf{e}$$

Where λ is a constant

We call **e** as an eigenvector and λ as eigenvalue of **C**

Example: Principal Components

6 spectral images
from an airborne
Scanner.
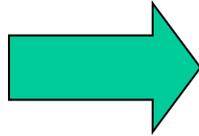
Channel	Wavelength band (microns)
1	0.40–0.44
2	0.62–0.66
3	0.66–0.72
4	0.80–1.00
5	1.00–1.40
6	2.00–2.60

Channel 1   Channel 2
Channel 3   Channel 4
Channel 5   Channel 6

(Images from Rafael C. Gonzalez and Richard E.
Wood, Digital Image Processing, 2nd Edition.
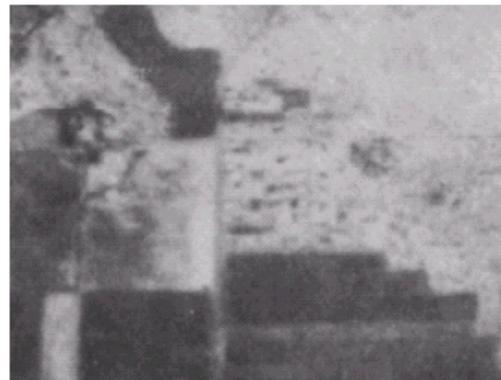
Example: Principal Components (cont.)

Component	λ
1	3210
2	931.4
3	118.5
4	83.88
5	64.00
6	13.40

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.

Example: Principal Components (cont.)

Channel 1 · Channel 2 · Channel 3 · Channel 4 · Channel 5 · Channel 6
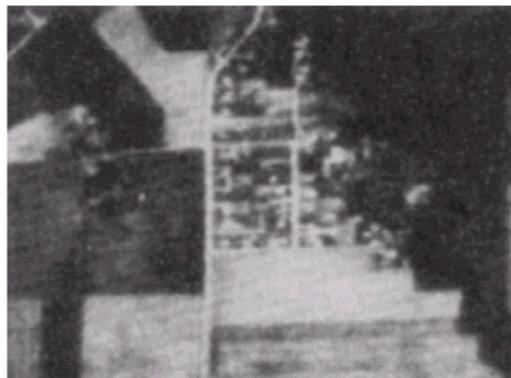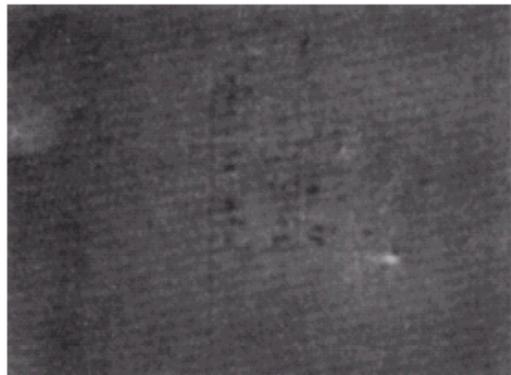
Component 1 · Component 2 · Component 3 · Component 4 · Component 5 · Component 6
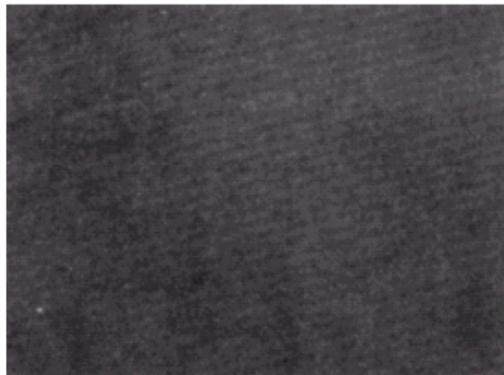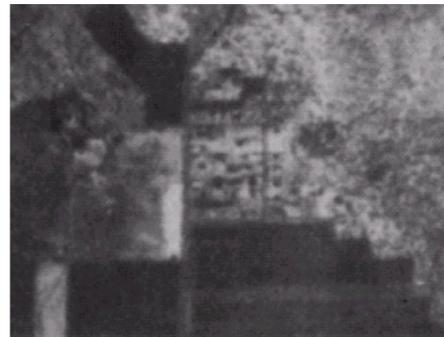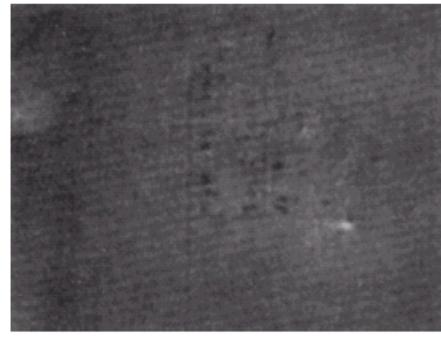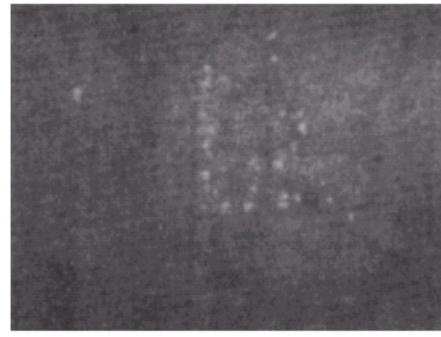
Original image

After Hotelling transform

Principal Components for Description
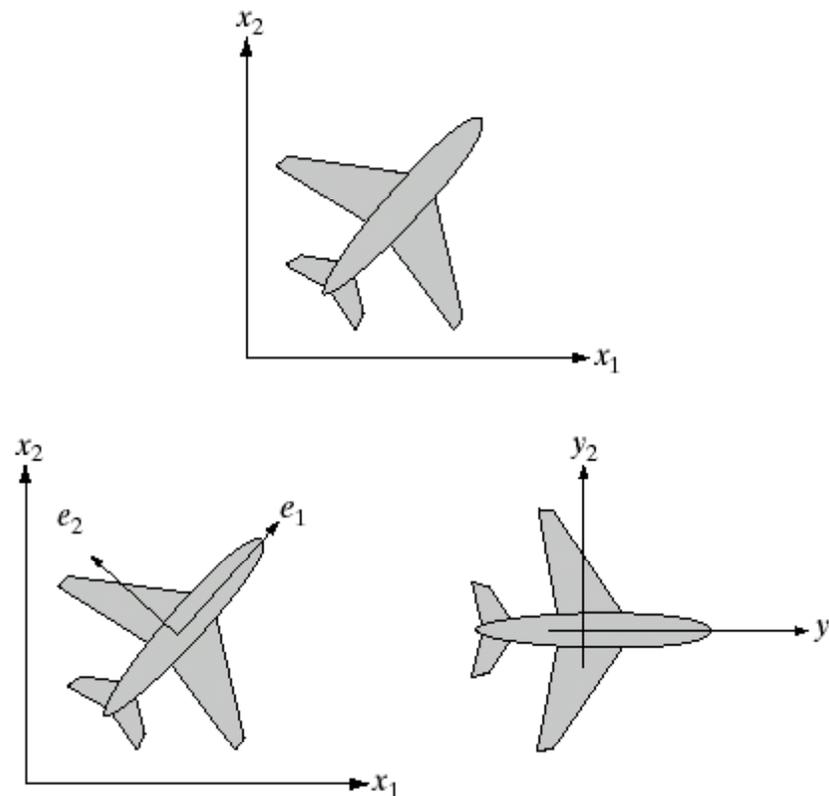
a
b c

FIGURE 11.29 (a) An object. (b) Eigenvectors. (c) Object rotated by using Eq. (11.4-6). The net effect is to align the object along its eigen axes.

(Images from Rafael C. Gonzalez and Richard E. Wood, Digital Image Processing, 2nd Edition.)

Relational Descriptors

a b

FIGURE 11.30
(a) A simple staircase structure.
(b) Coded structure.
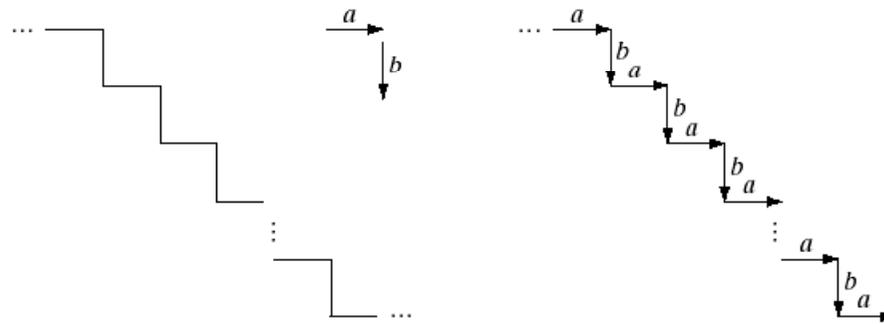
Relational Descriptors

FIGURE 11.31
Sample
derivations for
the rules $S \rightarrow aA$,
$A \rightarrow bS$, and
$A \rightarrow b$.
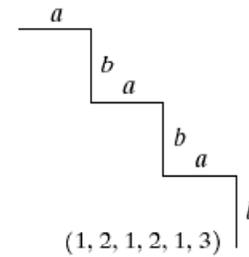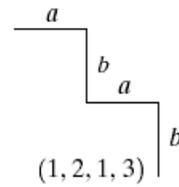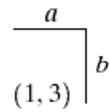
Relational Descriptors
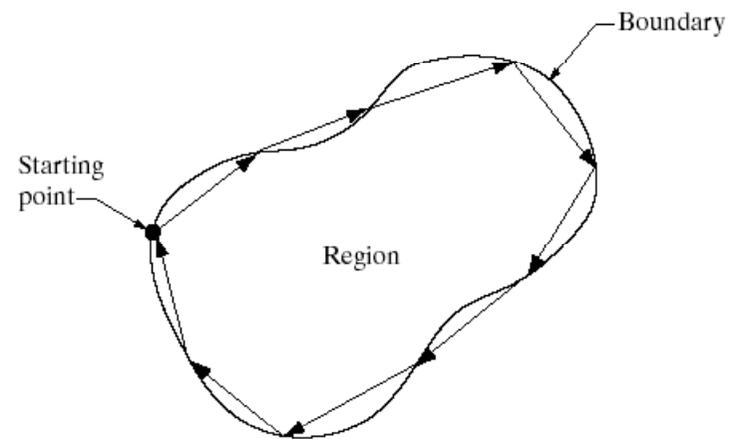
FIGURE 11.32
Coding a region
boundary with
directed line
segments.

(Images from Rafael C. Gonzalez and Richard E.
Wood, Digital Image Processing, 2nd Edition.

Relational Descriptors
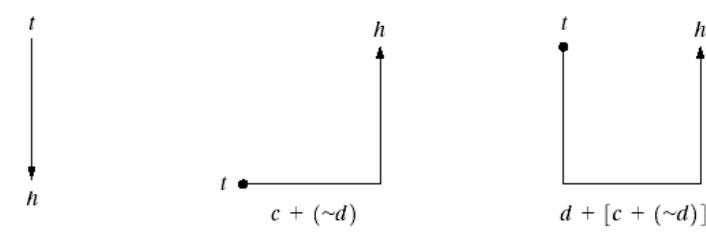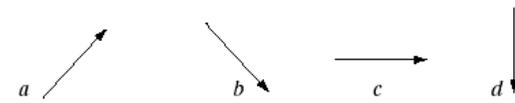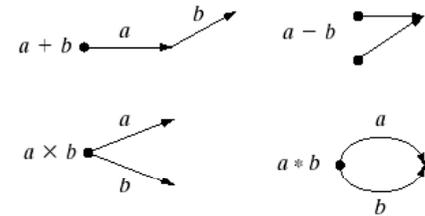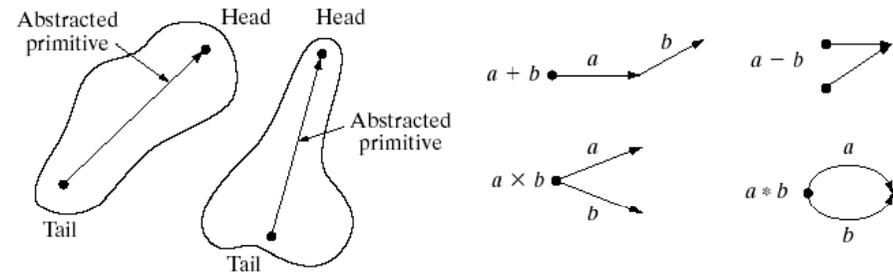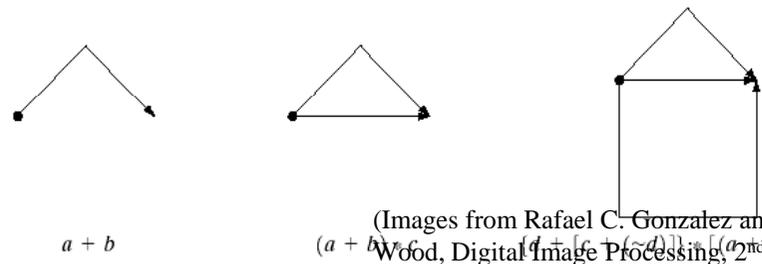
FIGURE 11.33 (a) Abstracted primitives. (b) Operations among primitives. (c) A set of specific primitives. (d) Steps in building a structure.

Relational Descriptors
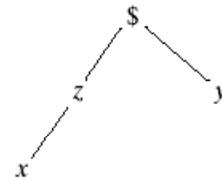
FIGURE 11.34 A simple tree with root $ and frontier *xy*.
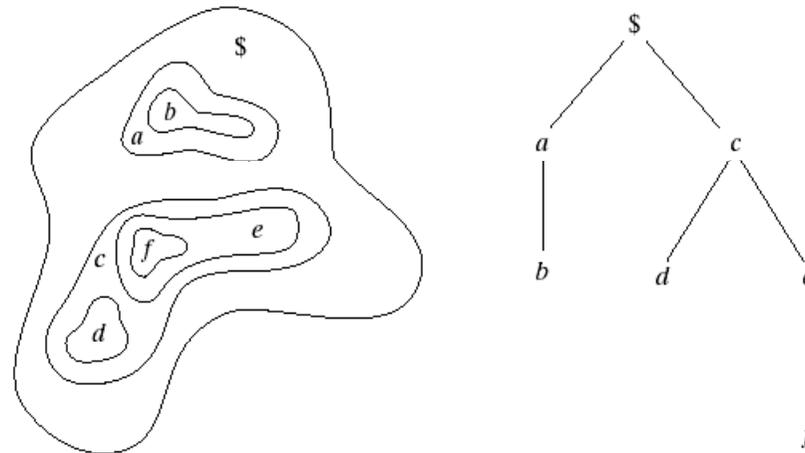
a b

FIGURE 11.35 (a) A simple composite region. (b) Tree representation obtained by using the relationship "inside of."

Structural Approach for Texture Descriptor

a
b
c

FIGURE 11.23
(a) Texture primitive.
(b) Pattern generated by the rule $S \rightarrow aS$.
(c) 2-D texture pattern generated by this and other rules.