IMAGE QUALITY

5

Image quality is a generic concept that applies to all types of images, regardless of their origin. For instance, it applies to medical images, photography, television images, and satellite reconnaissance images. In the context of radiology, the final measure of the quality of an image is the usefulness of the image in determining a meaningful and accurate diagnosis. More tangible and objective measures of image quality, however, are the subject of this chapter. Nevertheless, it is important to establish from the outset that the concepts of image quality discussed below are intrinsically related to the diagnostic utility of an image. Although it is true that huge masses can be seen on even low-quality images, and no amount of image fidelity will demonstrate pathology that is too small or faint to be detected, the true test of an imaging system and of a radiologist is the detection and accurate depiction of subtle abnormalities. With these types of diagnostic challenges, maintaining the highest image fidelity possible is crucial to the practicing radiologist. Consequently, understanding the features that comprise image quality is necessary in order to recognize problems and articulate their cause when they do occur.

As visual creatures, humans are quite adroit at visual cognition and description. Most people can look at an image and determine if it is "grainy" or "out of focus." One of the primary aims of this chapter is to introduce the reader to the terminology used to describe image quality and the parameters that are used to measure it. The principal components of image quality are *contrast, noise,* and *spatial resolution.* Each of these will be discussed in detail below.

CONTRAST

Contrast is the difference in signal between the background of the image and the structure one is looking for. For example, on a hand radiograph, a bone usually demonstrates high contrast, because the bone regions on the image are very "white" and the background is quite "dark." On a radiograph and most film images, the *signal* is related to the optical density of the image at some location on the image. When the radiograph is placed against a lightbox, the homogeneous light output of the lightbox

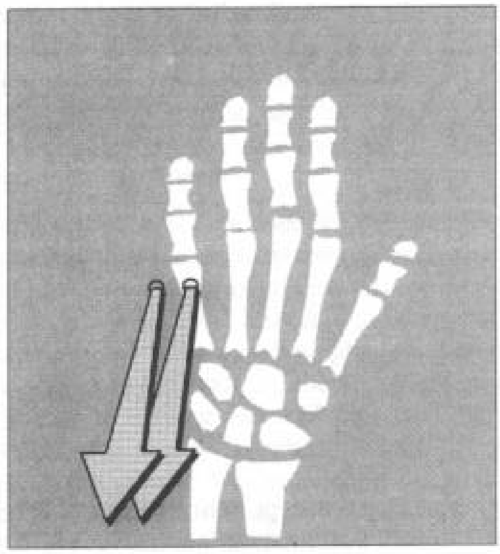

Figure 5.1. The radiographic contrast is simply the difference in optical densities between two adjacent areas.

is modulated by the distribution of optical density in the radiograph. The signal in this example is the number of photons being transmitted through a given area of the film (Fig. 5.1). The *signal difference* between the white areas under a phalanx and the dark areas behind the phalanges amounts to the contrast of the bone. In fact, the *radiographic contrast* is defined as the difference in optical density (OD) between a signal region and a background region on an image:

$$\text{Radiographic Contrast} = OD_A - OD_B$$

Radiographic contrast is the tangible result of acquiring an image of a patient or other subject. It is the *result* of both *subject contrast*, an intrinsic property relating to the interaction of x-rays with the patient's anatomy, and of the *detector contrast*, an intrinsic property of the detector system used to make the image.

Subject Contrast

Subject contrast relates to the ability of the *carrier wave* used in image formation to interact with structures of interest. For convenience, let us refer generically to the structure of interest as a lesion. It is understood, however, that in many instances it is the normal anatomy that is the structure of interest. The carrier wave is the type of energy that penetrates the patient and is used for image formation. In radiography and CT, the carrier wave is the x-ray; in ultrasound, it is sound waves; in MRI it is radiofrequency waves; and in nuclear imaging it is the gamma ray. The subject contrast is determined by the ability of the lesion to interact with the carrier wave differently from the surrounding tissue. The two components of subject contrast are the properties of the lesion itself and the characteristics of the carrier wave. The lesion has physical properties such as size, shape, composition, density, and location that determine its ability to interact with the carrier wave. The carrier wave has physical properties, such as photon energy, sound frequency or RF frequency, that influence its interaction with the physical properties of the lesion. The example of x-ray interaction with a lesion will be discussed below to illustrate the concept of subject contrast; however, the notion of subject contrast applies to all imaging modalities and carrier waves.

In general, the more that a lesion interacts with a carrier wave, the greater the subject contrast. In Figure 5.2, there is a lesion contained within a slab of solid tissue. As discussed above, the subject contrast of this lesion is related to physical characteristics of the lesion itself and of the x-rays (in this example). The probability of interaction between the x-rays and the lesion in this example is described by the *x-ray attenuation coefficient,* usually given the symbol μ. For a calcified lesion, the mass attenuation coefficient is higher than for a soft tissue lesion of the same size and density, thus demonstrating the effect that lesion *composition* has on subject contrast. Increasing density or size of the lesion will further increase its subject contrast. Because the attenuation coefficient is very dependent on the spectrum of energies in the x-ray beam used to produce the image, characteristics of the carrier wave can have a profound effect on subject contrast. In most cases there is very little that can be done to change the physical properties of the lesion. Therefore, efforts to maximize subject contrast usually focus on the properties of the x-ray beam. In mammography, for example, very low x-ray energies are used to maximize the subject contrast in the breast. This is effective because the attenuation coefficients are substantially higher at low x-ray energies, and consequently at low energies the important, albeit subtle,

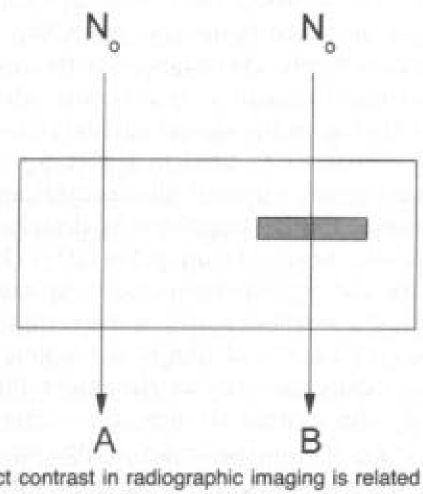

Figure 5.2. Subject contrast in radiographic imaging is related to the difference in x-ray transmission between two adjacent areas in an object. The intensity under a radiopaque object, *B*, is less than the x-ray intensity under a corresponding nearby area, *A*. The subject contrast in percent is calculated as: *Subject Contrast (%) = 100 × (A − B)/A (where A>B).*

differences in breast tissue demonstrate greater subject contrast, which is defined as:

Subject Contrast (%) = 100 × (A − B)/A (where A>B)

This definition of subject contrast (Fig. 5.2) is convenient because it ranges between 0 and 100%. Different definitions of subject contrast can be found in other textbooks, and the precise form is not too important. All definitions of subject contrast do embody the transmission of the carrier wave through the object relative to the background.

It is safe to say that, if an object has no subject contrast, it will not be seen on the image. However, for objects of interest that have very little subject contrast, the detector system can help amplify contrast. This is the topic of the next section.

Detector Contrast

Detectors in radiography can be screen-film cassettes (radiography), scintillation crystals coupled to light detectors such as photodiodes or photomultiplier tubes (digital radiography, nuclear medicine, CT), piezoelectric transducers (ultrasound), or antennas (MRI). All of these devices have the property that they interact with and record the presence of the carrier wave that has been transmitted through or has been emitted from the patient. Although the responsibility of the detector is straightforward, i.e., to detect, certain characteristics of the detector system can yield contrast enhancement or amplification right at the detector stage. Because of the profound difference between screen-film and electronic detectors, their characteristics will be described separately.

Screen-Film Systems as X-ray Detectors

The characteristic curve of a screen-film system is shown in Figure 5.3. Screen-film systems are discussed more fully in Chapter 6. Here, the characteristic curve is discussed in the context of how it affects contrast in the image. The characteristic curve describes the response of the detector system to different levels of x-ray exposure. The *response* in the case of screen film systems is the optical density (OD) of the film, and this is plotted as the ordinate (y axis) in Figure 5.3. There is obvious curvature in the characteristic curve for film, and this implies that screen-film systems are *nonlinear*. The equation that describes the linear portion of the characteristic curve in Figure 5.3 is OD α (Exposure)$^{\gamma}$. The γ is the "gamma" of the film screen system and is typically around 3.

The slope of the characteristic curve is the contrast of the detector. Because the characteristic curve of film is not a straight line, its slope varies, depending on the exposure that the screen-film system has received. Consequently, the contrast changes, depending on the exposure level. The contrast is plotted on Figure 5.3, along with the characteristic curve. It is clear that the contrast is highest in the midrange exposures, and toward the toe and the shoulder of the curve the contrast drops off significantly. This is consistent with clinical experience. On a chest radio-

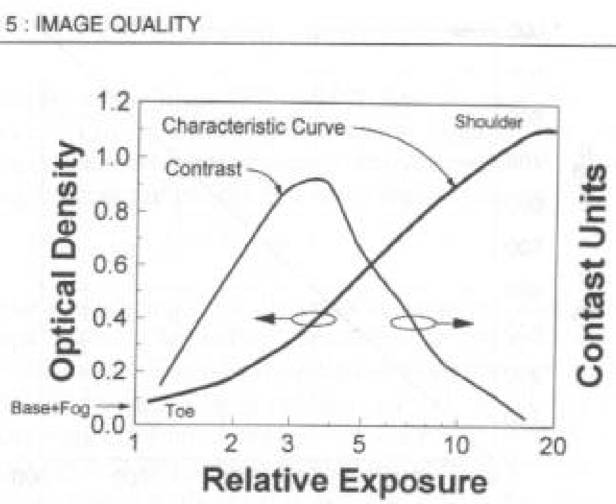

Figure 5.3. The characteristic curve for film is a graph of the optical density that results from a given exposure to light or to x-rays. The film will have a measurable OD even when not exposed, and this is the "base + fog" level. The contrast curve illustrates the films "γ," and is the slope of the characteristic curve. The contrast is lower in the underexposed toe region and in the overexposed shoulder region of the curve. Optimal exposures occur in the center, linear portion of the characteristic curve, where the contrast is highest.

graph, the ribs are easily visible in the lung region, where midlevel exposures typically occur. Under the mediastinum, however, the ribs are often seen with less clarity because the contrast in the detector is lower in underexposed areas.

Electronic Detectors

The nonlinear characteristic curve of screen-film detector systems is a consequence of the film's response characteristics. The amount of visible light emitted by the screen is a very linear function of x-ray exposure. Consequently, when a detector is composed of a phosphor material, such as an intensifying screen coupled to a light-sensitive electronic detector (e.g., a *photodiode*), the characteristic curve is usually very linear (Figure 5.4). In radiologic imaging, the output of most electronic detector systems is digitized and stored in a computer. Computed tomography is a good example of a detector system in which the signal from electronic detectors is digitized. Digital images have an interesting property with respect to contrast. Because digital image data is in a computer, the displayed characteristics of the image can be easily manipulated by the computer, as described in Chapter 3. In fact, the look-up tables used to display digital images can be used to increase image contrast significantly (Fig. 5.5). With digital images, contrast can be enhanced to the point at which the detector contrast is not a limitation of the contrast performance of the imaging system. Whereas screen-film systems are not capa-

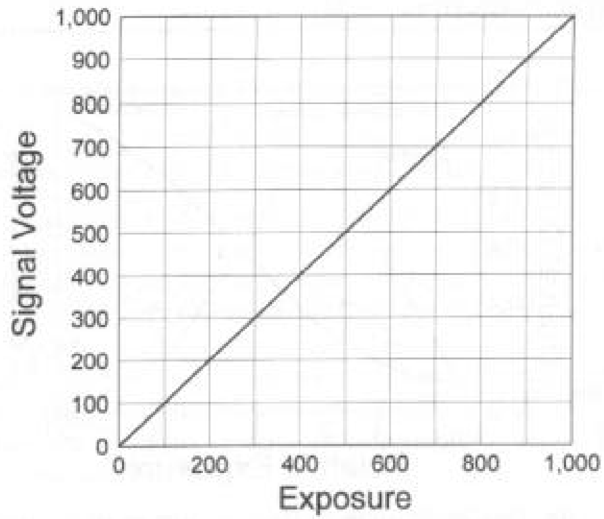

Figure 5.4. The characteristic curve for an electronic detector called a *photodiode*. The x-rays strike a phosphor material (such as a small piece of x-ray screen) and emit light, which is detected by the photodiode. The photodiode produces a small current (which is converted to a voltage) that is related to the intensity of light it detects. The characteristic curve for an electronic detector of this kind is very linear and has similar contrast throughout its entire range (i.e., the slope of this curve is a constant).

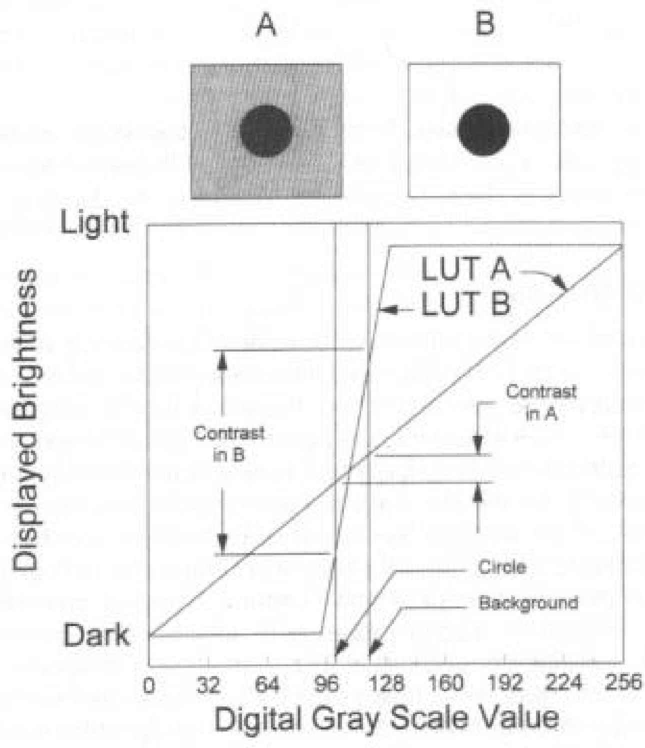

Figure 5.5. The displayed contrast of digital images can be enhanced by changing the shape of the look-up table (LUT) in the display system. For the same gray scale data stored in the computer, a narrower window (LUT B) will display the image data with greater contrast, as illustrated above by the high contrast dark circle on a white background (*B*). A wider window width (LUT A) demonstrates reduced contrast in the image (*A*).

ble of redisplay of the image with higher contrast, in digital systems the image contrast can be increased arbitrarily. Consequently, screen-film systems are considered *contrast-limited detectors*, whereas electronic/digital imaging systems are considered *noise-limited systems*.

NOISE

Imagine that it is raining on a tile patio. If we could count the number of raindrops that have landed on each tile in an hour, that number will vary from tile to tile. If it is raining hard, the number of raindrops per tile will be, on average, a large number. Lighter rainfall will result in a lower mean number of raindrops per tile. Light rain or hard, the number of drops per tile will vary slightly around the mean or average value. The most common way of analyzing this raindrop data would be to "bin" the data. For example, if there were 1000 tiles on the patio, and if the average number of raindrops per tile was 100, we would record the number of tiles that had between 0 and 10 drops, then record the number of tiles that had between 11 and 20 drops, and so on. This data is plotted in Figure 5.6. This type of plot is called a *histogram* and demonstrates the *distribution* about the mean. The shape and the lateral extent of the distribution is an excellent graphical description of the *noise distribution*. Figure 5.7 illustrates more realistic noise distributions, since there is *noise*

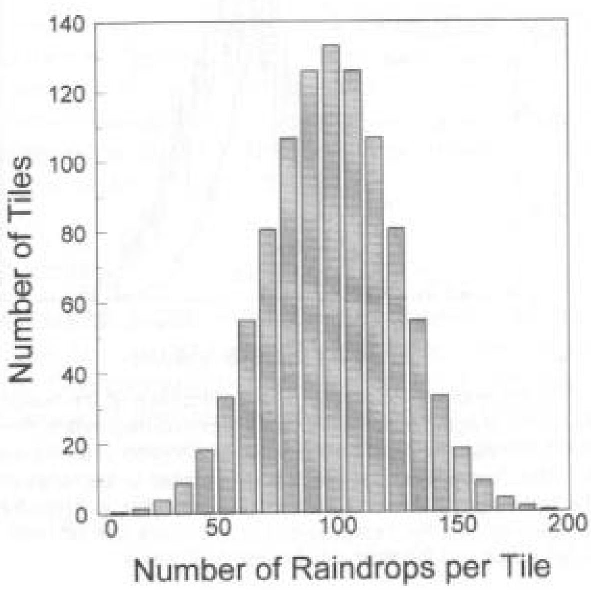

Figure 5.6. A histogram of the distribution of raindrop "intensity" that falls on some patio tiles. The x-axis corresponds to the number of raindrops that strike each tile, binned into categories of 0–10 drops, 11–20 drops, 21–30 drops, and so on. The y axis illustrates how many tiles on the patio could be found with the corresponding number of raindrops on them. This plot is a visual demonstration of the distribution of raindrops on the patio tiles. The total number of tiles (the area of the plot) is equal to 1000, because there are 1000 tiles on the patio. The mean number of raindrops is 100; however, some tiles experienced more and some tiles less raindrops.

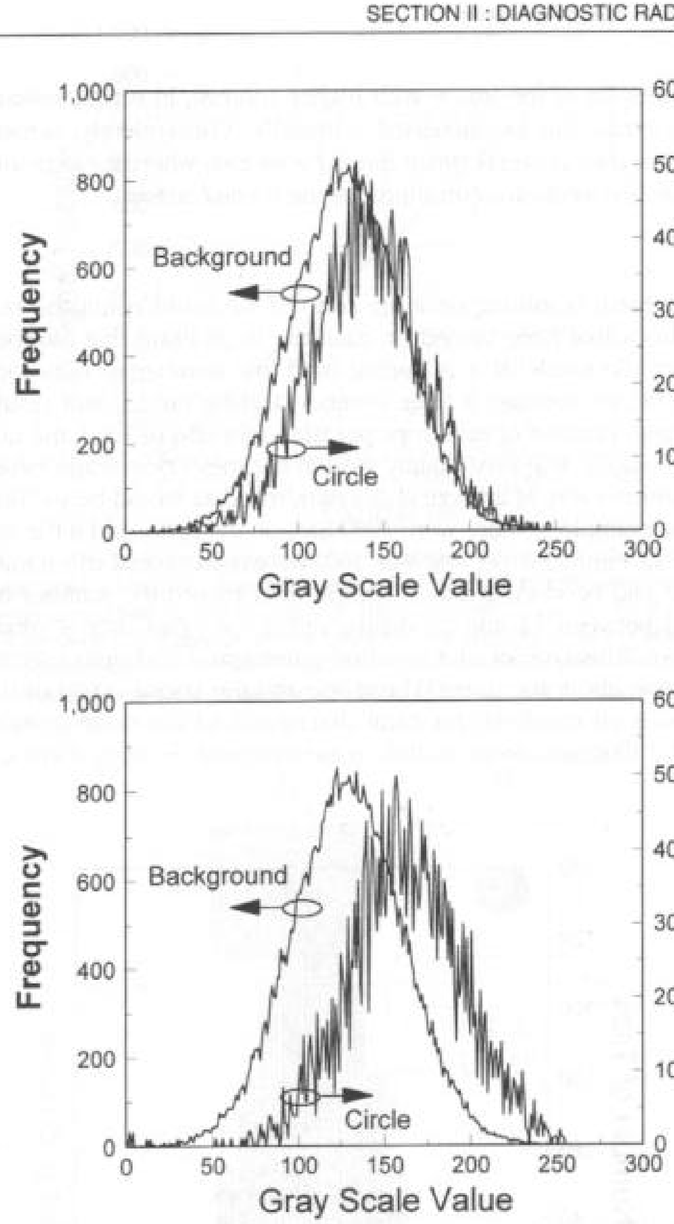

Figure 5.7. The top graph demonstrates a real distribution of the number of photons that strike pixels in an image. The left bell-shaped curve corresponds to the area outside a "white" dot on the image, the right-most curve is a histogram of the pixel values inside the "white" dot. The separation of the two curves is related to the contrast of the dot in the image. The bottom graph demonstrates a similar situation, in which the "white" dot has greater contrast against the background. For both plots, the left-most curve corresponds to the left y axis and the right curve to the right y axis.

in the noise distribution. The histograms in Figure 5.7 show the pixel counts in a signal region and in a background region. The lateral difference between the two curves in each plot is the result of contrast between the signal region and the background. There is more contrast (and the same amount of noise) in the bottom graph than in the top. The bottom graph shows a better signal-to-noise ratio, and consequently the

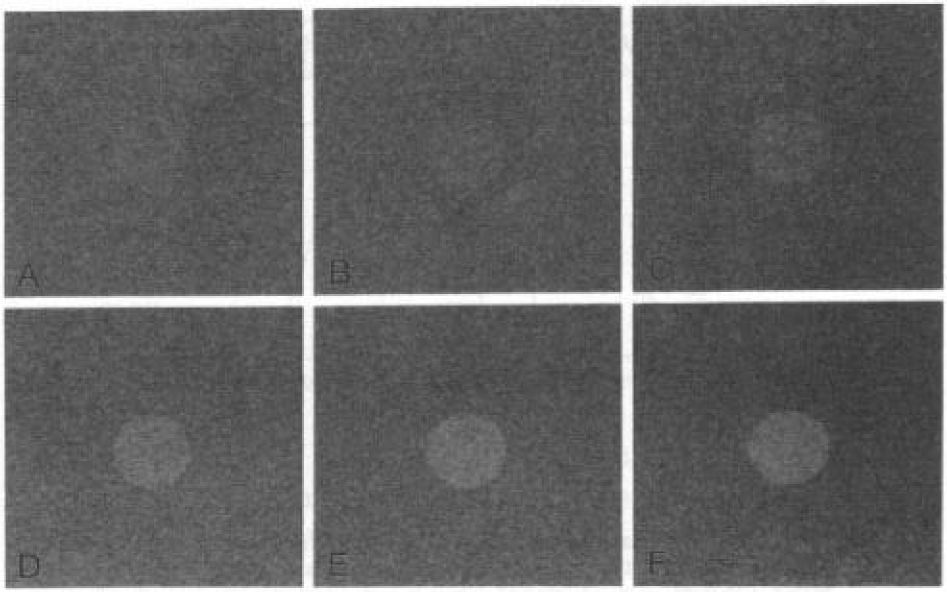

Figure 5.8. This is a series of six images of a white dot, each one showing different amounts of contrast. The noise in all images is the same, but the signal increases. Therefore, these images display a range of signal-to-noise ratios (SNRs). The histograms shown in Figure 5.7 are taken from these images; the top graphs correspond to panel A and the bottom graph to panel F.

signal region would be more apparent on an image. Visually, noise has the appearance of "graininess" or "static." An example of images with varying levels of noise is shown in Figure 5.8. The upper illustration in Figure 5.7 show histograms of the signal and background regions of the image shown in Figure 5.8*A*. The lower distributions in Figure 5.7 correspond to the image in Figure 5.8*F*.

Quantum Noise

In this section, the goal is to demonstrate the effect of noise on image quality. There are several commonly occurring noise distributions, for instance the Gaussian distribution (also known as the normal distribution), the binomial distribution, and the Poisson distribution. Each, of course, has its own mathematical description. Some observations can be made with respect to noise distributions. The formula that describes the Gaussian distribution is:

$$P(x) = \frac{1}{2\pi} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$$

The important observation that can be made about the Gaussian distribution is that there are *two parameters* that determine its shape, the mean value, μ, and the standard deviation, σ. The value of x is the *dependent variable,* not a parameter that affects the shape of the distribution. These parameters are illustrated in Figure 5.9. The flexibility of having two parameters, one that describes the average value and the other that

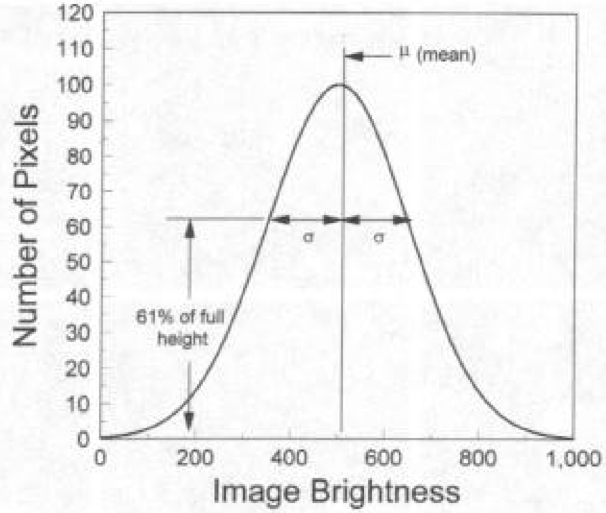

Figure 5.9. The Gaussian distribution, also called the normal distribution, is shown. Because the distribution is symmetrical, the mean value of the distribution is in the center. The width of the distribution, which is a measure of the noise in the distribution, is given by the value of σ. The value of σ can be determined by measuring the width of the curve at 61% of its total height. This width is equal to 2σ.

describes the noise levels is quite useful, and the Gaussian distribution is by far the most widely used statistical distribution in scientific measurement. An example of where this would be useful is for measuring the height of children in a classroom. In this example, the mean value is not necessarily related to the noise, and therefore it is useful to have the two parameters, one describing the mean height and another describing the fluctuations in height about the mean value.

When counting a large number of individual events or *quanta,* such as the number of raindrops striking patio tiles or the number of x-ray photons per pixel in an image, the Poisson distribution is used. The Poisson distribution has a completely different formal mathematical expression than that for the Gaussian distribution, but there is only one parameter that describes both the mean and the standard deviation. To a first approximation, the Gaussian distribution can be used to represent the Poisson distribution, except that *the mean is equal to the standard deviation squared,* $\mu = \sigma^2$, or equivalently, $\sqrt{\mu} = \sigma$. The square of the standard deviation, σ^2, is called the *variance*. The Poisson distribution is easy to work with for the following reason: The mean is very easy to measure. For example, one just counts the total number of raindrops hitting the patio and divides by the number of tiles, and the result is the mean number of raindrops per tile. Let's call this mean N. Because this phenomenon is governed by the Poisson distribution, the noise is equal to the square root of the mean, $\sigma = \sqrt{N}$. Consequently, no other measurements are needed to estimate the noise. Because the standard deviation due to *quantum noise,* σ, is a function of the mean, it is easy to predict from the mean exposure what the quantum noise will be in an image. *As the mean number of photons (N) used in making an image increases,*

Table 5.1
The Relationship Between the Mean Number of Photons Used in the Image, the Standard Deviation in the Mean, and the Percent Noise Levels in the Image[a]

Mean Number of Photons, N	Standard Deviation, $\sigma = \sqrt{N}$	Noise, $100 \times \sigma/N$ (%)
100	10	10
500	22	4.5
1,000	32	3.2
5,000	71	1.4
10,000	100	1.0
50,000	224	0.4
100,000	316	0.3

[a]The noise *perceived* in the image, i.e., the percent noise, decreases as the mean number increases.

the standard deviation, σ, also goes up. However, and *very importantly,* σ does not increase proportionately with the mean, but rather as the square root of the mean. Said differently, the noise σ increases *more slowly* than the mean. It is the *noise level relative to the mean* (also called the *fractional noise*), σ/N, that is visually perceived in an image. Consequently, as the mean number of x-ray photons (N) that are used in forming an image increases, the fractional noise decreases. This principle is demonstrated in Table 5.1. In x-ray imaging, the amount of x-ray exposure can be increased to make an image with better image quality, but the patient's dose will increase as well. This is the all-important compromise in radiographic imaging.

Quantum noise refers to the noise due to the influence of the *number* of individual particles or photons *(quanta)* used in the formation of the image. Quantum noise usually obeys the Poisson distribution, and consequently we know *a priori* that as the number of photons increases, the perceived noise in the image will decrease. *In well-designed x-ray systems, the quantum noise should be governed by the number of x-ray photons (quanta) absorbed in the detector.* Nevertheless, it is possible to have systems that are not x-ray quantum limited. For example, if a video camera is pointed at a fluoroscopic screen (shown in Fig. 5.10), a radiographic image can be made. However, although a large number of light photons are emitted from the screen when one x-ray photon is detected (about 200), only a very small fraction of these light photons ever get to the video camera (about one photon in 1000). Consequently, for every 5 x-ray photons absorbed in the screen, only 1 light photon will be recorded by the video camera. In this example, there are fewer light quanta than x-ray quanta used in forming the image, and therefore the quantum noise will be governed by the number of light quanta, not the number of x-ray quanta. Because the number of light quanta is less than the number of x-rays, the quantum noise contribution from this stage in the detector will be greater and will dominate the noise in the system. When an x-ray system is not noise limited by the number of x-ray quanta, but rather by a *secondary quantum sink,* x-ray quanta are wasted (not recorded). This means that the patient is exposed to radiation that is not

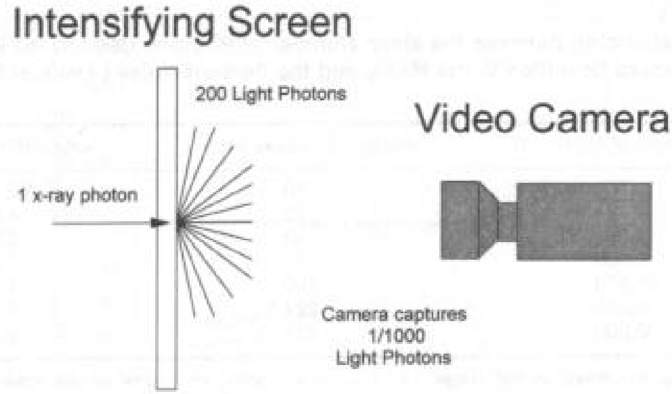

Figure 5.10. An example in which the quantum noise of the imaging system is not determined by the number of detected x-ray photons, but instead by the number of optical photons. One x-ray photon, on average, results in the emission of 200 light photons from the screen. The poor optical coupling in the system results in the collection of only 1 in 1000 of the light photons by the video camera. This means that, on the average, five x-ray photons must be detected to record one light photon in the video camera. Such a system is an example of wasted x-ray exposure to the patient. Because the number of light photons recorded is less than the number of x-ray photons detected, the fractional noise of the light photon distribution is greater than that of the x-ray photon distribution. Therefore, the noise in the system is dominated by the optical photon noise, not the x-ray quantum noise, as it should be. This is an example of a secondary quantum sink.

recorded, and that means that much of the radiation dose to the patient is wasted. This is a very undesirable situation; *all properly designed x-ray systems should be x-ray quantum limited.*

Other Sources of Noise

Although x-ray quantum noise is and should be a limiting factor in x-ray detector design, there are other sources of noise that affect radiologic images. In screen-film systems, the granularity of the film emulsion is a source of random noise (random noise is also called *stochastic noise*). The manufacture of radiographic screens sometimes results in small variations in the thickness of the screen across its surface, and this can lead to a slight difference in film darkening from one region to another. This type of noise is a source of artifact that will show up repeatedly on all images that are produced with the affected intensifying screen. Because this is a reproducible type of noise, it is not random in nature and therefore is called *nonstochastic noise*. It is also called *structured noise*. Structured noise can appear as a result of the manufacturing process in virtually all imaging systems. Image intensifiers, video cameras, scintillation cameras, and even computed tomography (CT) scanner detector systems can experience structured noise problems. A dried splash of iodinated contrast media on the input surface of an image intensifier will cause structured noise.

Imaging systems that use electronics often suffer from electronic noise, especially when high amplification is used. Often, with very large patients, a fluoroscopic system will reach the point at which no more x-ray photons can be emitted from the x-ray tube because the automatic brightness control is adjusted to not exceed a certain maximal exposure rate (10 R/min is the legal maximum, adjustment of fluoroscopic systems will vary). To maintain image brightness after the maximal exposure rate is achieved, the automatic gain control of the video camera will become operative, and will amplify the video signal so the image on the monitor appears brighter. Because in this situation the x-ray quantum noise is already very high (since the image is photon-starved), amplifying the noisy image amplifies the noise as well as the signal.

SIGNAL-TO-NOISE RATIO

The concepts of contrast and of noise were discussed above. Contrast is essentially *the signal* in diagnostic imaging, something that we strive to maximize. Noise is something that we try to minimize in the design and operation of imaging devices. There is a convenient way to combine the notions of contrast and of noise into a single parameter, the signal-to-noise ratio (SNR). The contrast to noise ratio (CNR) is also used:

$$SNR = \frac{Signal}{Noise} \qquad CNR = \frac{Contrast}{Noise}$$

The SNR sounds like, and is, an engineering term commonly used in other fields. It is a convenient way of expressing the signal levels relative to the noise levels. Because it is desirable to maximize the signal and minimize the noise, we would like the SNR to be very high (Fig. 5.11).

The SNR is the holy grail when it comes to *contrast resolution*, that is, the ability to see large but faint objects in an image. A clinical example of the importance of contrast resolution is in the detection of liver metastases on a CT scan. The often subtle difference between the density of the metastatic lesions and the density of the surrounding normal liver parenchyma needs to be large enough to be visible against the noisy background. If the contrast is too subtle, or the noise is too great, the lesion will not be visualized. The ability to perceive large area but very low contrast lesions is directly (not necessarily linearly) related to the SNR. As the SNR for an imaging system goes up, the better that imaging system is at producing images capable of good contrast resolution. Consequently, the SNR is one very good measure of the performance of an imaging system. It is one of the most important concepts regarding image quality. The SNR will be further explored later in this chapter.

SPATIAL RESOLUTION

Spatial resolution is a measure of how good a system is in producing images of very small objects. The classic definition of spatial resolution relates to the distance that two objects must be separated so that they be distinguished as two objects instead of one. The *limiting spatial resolution*, that is, how small of an object can be seen with an imaging system, is but one useful description of the resolution capabilities of an imaging system. Before completely describing the spatial resolution capabilities of

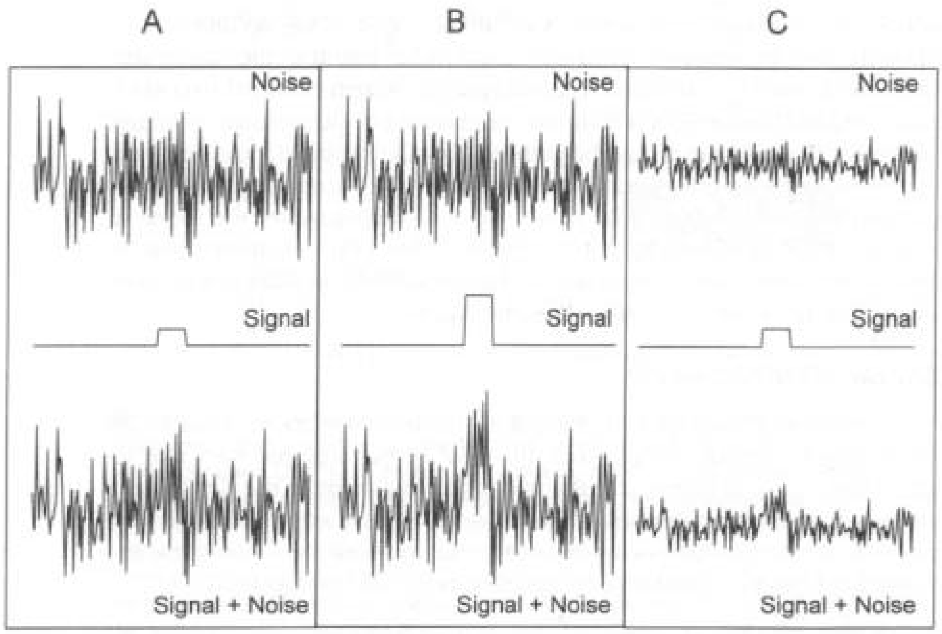

Figure 5.11. The effect of the signal-to-noise ratio (SNR) on the ability to detect a signal. **A.** The signal is small compared with the noise and therefore, when the signal is superimposed on the noise (as is the case with medical images), it is impossible to determine with any certainty whether the signal is there. **B.** The signal is larger, and therefore a larger SNR occurs. Here, the signal can be seen in the noisy background. **C.** The signal is the same height as in A; however, the noise is made smaller. Here again, the SNR is improved over that in A, and the signal can be seen in the noisy background. When the noise in an image is reduced, smaller signal amplitudes become detectable. Low noise images therefore exhibit better contrast resolution.

an imaging system, including a discussion of how well the system resolves larger objects, it is useful to discuss why resolution is lost in imaging systems.

Physical Mechanisms

There are a myriad of physical mechanisms that are responsible for the loss of resolution in imaging systems. Here, one basic example will be used to demonstrate how physical processes involved in imaging an object or patient give rise to a loss of resolution. The example of the resolution loss in a radiographic intensifying screen will be discussed. (Intensifying screens are discussed in detail in the next chapter.)

The purpose of the intensifying screen is to absorb x-rays that have passed through the patient and, in doing so, they emit light that then exposes the radiographic film which is in tight contact with the screen. It is desirable to absorb as many x-ray photons as possible, and so from this perspective thicker screens would be better. However, thicker screens permit the light that is emitted during an absorption event to spread a further distance, resulting in a greater blurring of the light pattern. The thicker screen, therefore, results in a greater loss of resolution than the thinner screen. This is illustrated in Figure 5.12, which shows the classic compromise that needs be made in designing an imaging system. Thicker

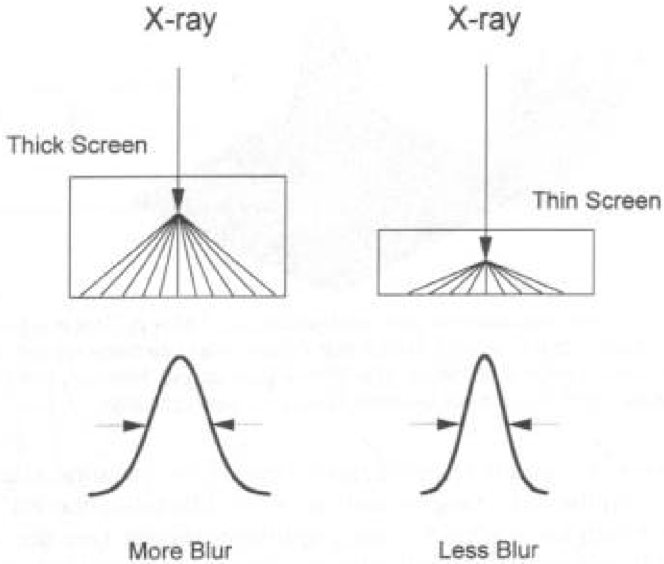

Figure 5.12. A physical example of how an imaging system can degrade the spatial resolution of an image during the recording process. The thicker screen detects a greater fraction of the x-rays that are incident upon it; however, the blur resulting from the diffusion of light within the screen is greater. The thick screen has more blur and poorer spatial resolution than the thin screen.

screens are more sensitive, and that is good; however, this sensitivity is achieved by losing some spatial resolution. The thinner screen does not absorb as many x-ray photons, so it is less sensitive than the thicker screen, but it has better resolution.

The original x-ray photon is essentially a perfect point input to the imaging system, and the spread of light causes a *blurred* version of that sharp input. The distribution and magnitude of light spread can be measured for a given imaging system. For systems that use other forms of energy other than light, the blurring of the signal in whatever form can be characterized as well.

Characterizing Spatial Resolution in the Spatial Domain

In the intensifying screen example given above, a perfect point input of energy in the form of a narrow beam of x-ray photons is absorbed, and the resultant visible light spreads out before exposing the film. The blur pattern itself is a good description of the resolution capabilities of the system. The response of a point input to an imaging system is called the *point spread function* (PSF) and describes the lateral spread of information in the imaging system. An example of the PSF for the light spread in a screen is illustrated in Figure 5.13. The spread of light occurs in both the x and the y directions, and therefore the PSF is often expressed as a function of these dimensions, PSF(x,y).

One of the problems in measuring the PSF of an imaging system is that it is very difficult to collimate a beam of x-ray photons into a perfect "point." A perfectly small point input would require a hole to be drilled

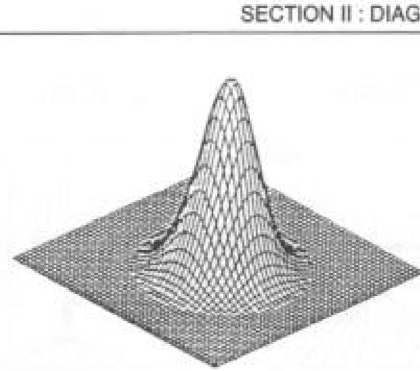

Figure 5.13. Isometric plot: the point spread function, PSF(x,y). This is a two-dimensional distribution of the amount of blur that occurs in all directions around the point where the x-ray photon is absorbed. This PSF is symmetrical; however, some imaging systems have PSFs that are not symmetrical (e.g., a video camera).

into a sheet of lead and to be so small that very few photons would ever pass through the hole. Despite such technical difficulties, the PSF can be measured with some effort by using a *pinhole camera*. One way to ease the problem of getting enough x-ray photons through a pinhole is to use a very narrow slit cut into the lead sheet. Instead of a point input to the imaging system, the lead slit collimates the x-ray beam into a line input to the imaging system. The response of the imaging system to such a line input is called the *line spread function* (LSF). If a line input is directed towards an imaging system along the y axis, the LSF is measured perpendicular to it, along the x axis. Because the placement of the axes is arbitrary, the LSF is usually described as a function of the distance along the x axis, LSF(x).

In some instances, it is technically difficult to make an appropriate exposure by using a lead slit. If a pinhole or lead slit is not available, sometimes just the edge of a sheet of lead is used. This allows a much larger area of the image receptor to receive radiation, and the imaging system's response to an edge input is called the *edge spread function* (ESF). Like the LSF, it is usually defined as a function of the x dimension, ESF(x). The PSF, LSF, and ESF are illustrated in Figure 5.14. There is a mathematical relationship between these various spread functions, and although the details of the mathematics are beyond the scope of the present discussion, the conceptual relationship is worth exploring in further detail. The mathematical operation that allows one to calculate the LSF from the PSF is called the *convolution*. Convolution is a recurrent theme in medical imaging, and therefore a brief, graphical explanation is in order.

The PSF describes the imaging system's response to an infinitely narrow spike of radiation. An infinitely narrow (spike-like) input to the imaging system is called a *delta function* (δ-function). When a slit of radiation exposes an imaging receptor, the LSF result is equivalent to a large number of PSF responses, all lined up in a row. As can be seen in Figure 5.15, to calculate the LSF from a line of PSFs, the contribution from each PSF in the line must be summed at each point in the image. This summation or integration is the convolution operation. This convolution business can be taken even a step further. Imagine that a hypothetical

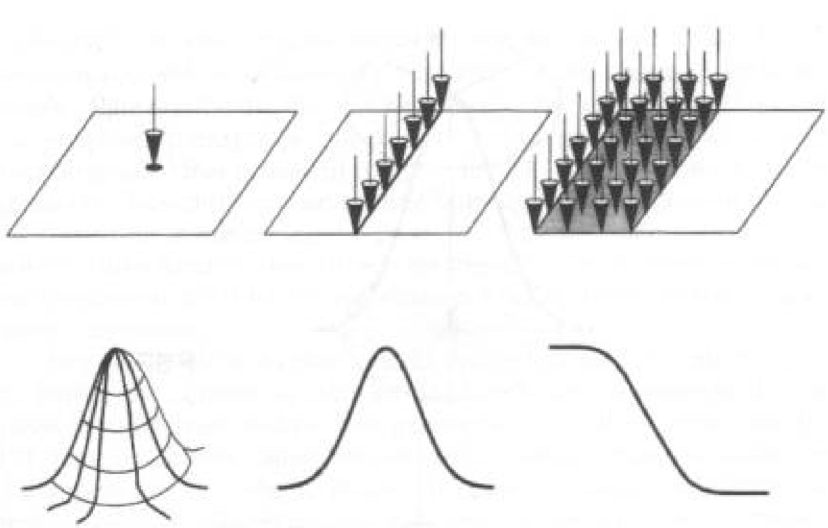

Figure 5.14. The differences among the point spread function [PSF(x,y)] (*A*), the line spread function [LSF(x)] (*B*), and the edge spread function [ESF(x)] (*C*). The PSF is measured by putting an infinitely narrow point into the imaging system. For an x-ray system, for example, a very small hole in a sheet of lead can be used to generate a point input. To measure a LSF for an x-ray system, a narrow slit in a sheet of lead is imaged, and the profile of the resulting line on the image is measured perpendicular to the line. This profile is the LSF(x). If just an edge of the sheet of lead is imaged, a profile across the edge perpendicular to it is the ESF(x).

imaging system capable of *perfect resolution* were used to image a patient's chest anatomy. Let us describe this perfect image symbolically as $I_{perfect}(x,y)$. Using this perfect image, the image that would be produced by a real imaging system, $I_{real}(x,y)$, can be calculated by convolving the point spread function of the imaging system, PSF(x,y), with the perfect image:

$$I_{real}(x,y) = I_{perfect}(x,y) \otimes \otimes PSF(x,y)$$

The symbol $\otimes$ represents the convolution operation, just as the symbol + represents addition. Two of these symbols are used above because the convolution occurs in two dimensions, the x and the y dimensions. The reader may wonder why one would be interested in degrading a perfect image into one poorer in resolution. A mathematical understanding of how an imaging system degrades the image, of course, is the first step in correcting the problem. Because it is possible to degrade a perfectly good image by convolving it with the PSF of an imaging system, might it be possible to convolve the real image with some function that *undoes* the PSF, so that one gets the perfect image out? Yes, in fact, the convolution operation does allow the restoration of images that have been blurred by an imaging system. This type of restoration is performed routinely in satellite reconnaissance imagery, where the distortion effects of the atmosphere are corrected. The Hubble telescope, a satellite-borne telescope launched by the space shuttle in 1990, made news when it was

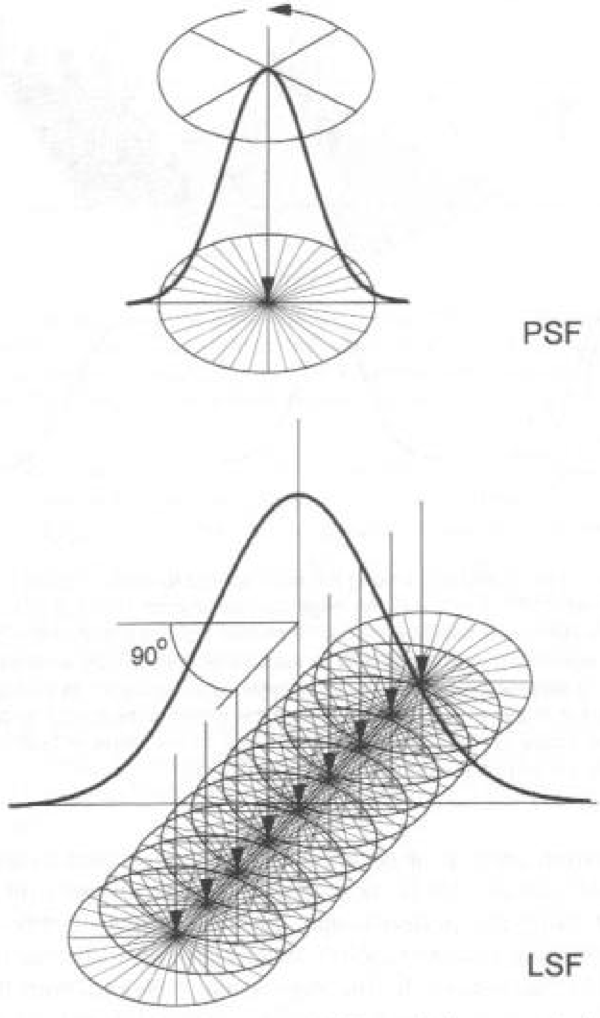

Figure 5.15. The LSF is really a series of PSFs overlaid onto each other along the slit. In mathematical terms, the LSF can be computed by convolving the PSF with the slit.

discovered that its designers had miscalculated and the telescope lens was not focused correctly. However, the PSF (the blur function caused by the lack of proper focus) for the Hubble telescope was measured by pointing the telescope at a very distant, solitary star whose input was essentially a δ-function. From the measured PSF, a function was derived that, when convolved with the out-of-focus images that the telescope produces, results in perfectly focused images. This restoration process is called *deconvolution*. The process, however, comes at a price—improving the spatial resolution by deconvolution will increase noise in the image.

Because convolution or deconvolution is a mathematical procedure, it can be conveniently performed on images that are stored in a computer, i.e. on digital images. All modern satellite images are produced in digital form, including those produced by the Hubble telescope. Film radiographs, however, are not digital, so it is not possible to use deconvolution routinely to enhance the resolution of the images. A good example of the routine use of convolution does exist in radiology, however. In computed tomography, the reconstruction process (called simple backprojection) that is used to compute the CT image from the acquired projection data causes a characteristic blurring of the image information. To correct for this blurring, the projection data are mathematically convolved with a function that corrects for the blur. The "filtered" in *filtered backprojection* refers to the mathematical filtering of the projection data using convolution.

In a world where medical images are increasingly becoming digital in format, the concept of convolution is necessary to understand. It is possible to convolve images with mathematical functions other than the PSF. Other functions, called *kernels,* can be used arbitrarily to influence the appearance of a digital image. For example, image smoothing (to reduce noise) or edge-enhancement (to increase the appearance of edges) can be performed using convolution. The convolution operation is the basis for much of what is generically called *image processing.* Computers perform convolution mathematically, but physical processes such as the blurring of the x-ray input by the spread of the light in the screen perform convolution *optically.* Deconvolution is performed optically when one puts on prescription eyeglasses. The optometrist's prescription essentially defines the (optical) deconvolution kernel (i.e., the curvature of the eyeglass lens) that corrects for the blur imposed by an improperly shaped lens in the eye.

The Frequency Domain

The previous section dealt with spatial resolution, and parameters that describe resolution. The expressions $PSF(x,y)$, $LSF(x)$, and $ESF(x)$ indicate functions of spatial distance, meaning that the units of x or y are units of *length* (e.g., millimeters, centimeters, or inches). In the parlance of image science, these functions describe various blurring phenomena in the *spatial domain.* There is another very useful way of describing resolution mathematically, using functions that are in the *frequency domain.* Before proceeding with a description of frequency domain functions, it may be useful to lay a foundation for understanding by discussing the frequency domain from a graphic and conceptual, rather than mathematical, point of view.

The sound waves that strike our ear drums are disturbances that change the pressure of the air next to the ear drum as a function of time (they are in the *temporal domain*). One way to represent this would be by the expression $P(t)$, where P is the pressure measured right near the eardrum and t is time. A plot of such a function is shown in Figure 5.16. When the peaks and valleys of a sinusoidal sound wave are *separated by shorter periods of time,* we hear these as *higher frequencies.* The middle A key on a piano is tuned to emit precisely 440 sinusoidal waves per second. The units of temporal frequency are *inverse seconds,* 1/sec, also called Hertz (Hz). Middle A corresponds to 440 Hz. Sound waves whose

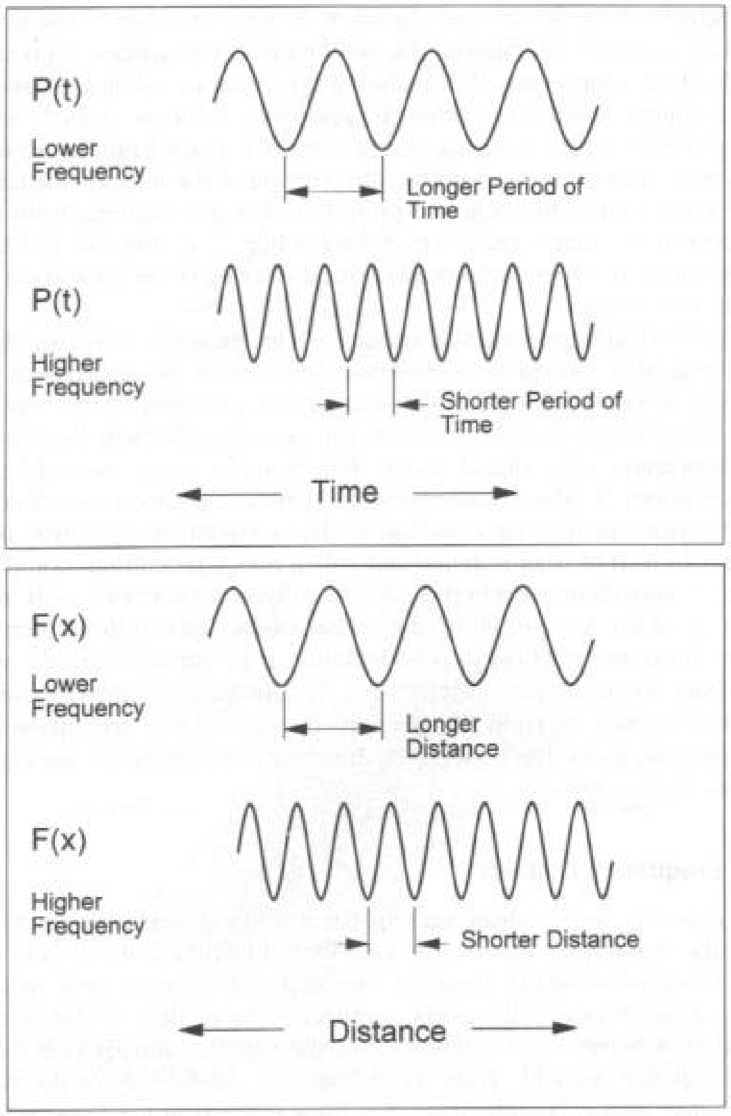

Figure 5.16. The notion of temporal frequency is familiar to all of us with good hearing. The notion of spatial frequency is less familiar but is very similar in concept to temporal frequency. Larger objects with greater separation between the objects correspond to low spatial frequencies. Small objects that are close together (e.g., microcalcifications on a mammogram) exhibit higher spatial frequencies.

peaks are *separated by longer periods* of time are *lower frequency* sound waves. There is an analogy between temporal frequency, measured in inverse time (1/sec), and spatial frequency, measured in inverse distance (1/mm). Just as high temporal frequencies correspond to sound waves whose peaks are separated by shorter periods of time, high spatial frequencies correspond to objects separated by shorter distances. High spatial frequencies describe objects that are small and physically close together, such as a cluster of microcalcifications on a mammogram. Low spatial frequencies correspond to large objects that are separated by large distances. The very wide, smooth changes in optical density that the mediastinum imposes on a chest radiograph are low in spatial frequency.

The Frequency Domain Description of Spatial Resolution

The *modulation transfer function* (MTF) is a graphical description of the resolution capabilities of an imaging system. A typical MTF is shown in Figure 5.17. A separate look at each axis will clarify the concept. An object, for example a gallstone, that has X amount of contrast is imaged by a system. If it retains 80% of that contrast (0.8 X) on the resultant image, the y value on the MTF graph would be 80%. Because contrast ranges from 0 and 100%, the y axis of the MTF plot usually spans this range as well. The x axis of the MTF graph corresponds to the size of the object that is being measured. The numbers to the left of the graph (low spatial frequencies) represent big objects. The gallstone was fairly big, and its size is indicated on the graph. Smaller objects are described by the higher spatial frequencies toward the right side of the graph. Because most imaging systems are better at retaining the contrast of large objects than of small ones, the MTF typically is good (close to 100%) at low spatial fre-

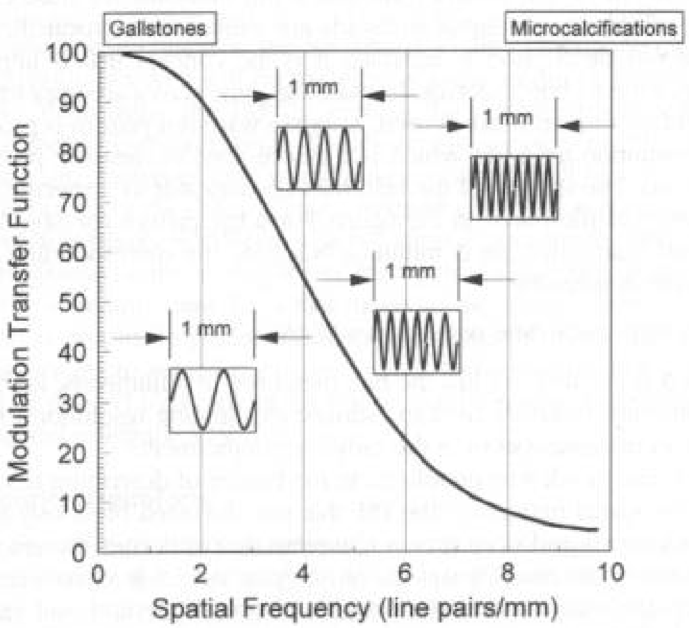

Figure 5.17. The modulation transfer function (MTF) is a plot of the contrast transfer properties of an imaging system (on the y axis) as a function of the size of the object (on the x axis). The size of the objects (represented by sinusoidal waves) are shown as insets for spatial frequencies of 2, 4, 6, and 8 line pairs per mm. An imaging system reduces the contrast of the imaged object when some of the information from peaks bleeds into the valleys, and vice versa. When the peaks and valleys are closer together (at higher spatial frequencies), this creep of information from one area to another results in a more severe reduction in contrast. Consequently, the MTF typically gets worse at higher spatial frequencies.

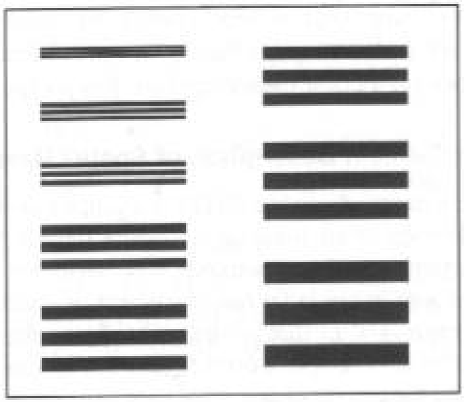

Figure 5.18. Line pair phantom. Phantoms such as this are used to measure the resolution capabilities of an imaging system in the field. This "quick and dirty" kind of measurement provides less information than an MTF, but it is a very practical way to monitor the resolution of an imaging system on a routine basis.

quencies (large objects) and gets smaller at the higher spatial frequencies (small objects).

The details of the MTF are such that the "objects" it describes are sinusoidal, but this is a mere technicality. Big sinusoids are close enough to big gallstones, and small sinusoids are similar to microcalcifications, so that our description is accurate. It is the concept that is important here, not the technical details. In radiology, sine waves are impractical to use, and so a square wave is used. A square wave test pattern is produced by a resolution *template,* which is a device used to measure resolution (Fig. 5.18). The spacing of the lead bars corresponds to a specific spatial frequency, as illustrated in the figure. For a bar pattern (or object) with bars and spaces that are Δ millimeters across, the corresponding spatial frequency is given by:

$$\text{Spatial Frequency (line pairs/mm)} = 1/(2\Delta)$$

where Δ is the width of just the bar, measured in millimeters. Resolution templates are routinely used to estimate the limiting resolution capabilities of an imaging system in the radiology department.

It is fair to ask why people go to the bother of describing resolution using the spatial frequency. The PSF that was discussed previously is fairly straightforward and does give a complete description of the resolution properties of an imaging system. An imaging system is often a series of devices; the signal enters the imaging chain at one end and cascades through the various components in the imaging chain until an image is produced at the end. A fluoroscopic system is a good example of a multicomponent system (Chapter 7), where the x-ray signal cascades through the input phosphor, the photocathode, the electron optics, the output phosphor, the optical lens, the video camera, and the electronics of the camera, before the image reaches the TV monitor. Those who design imaging systems measure the resolution characteristics of each component in the imaging chain. It is the weakest link in the imaging chain that determines the resolution capabilities of the entire imaging system. Iden-

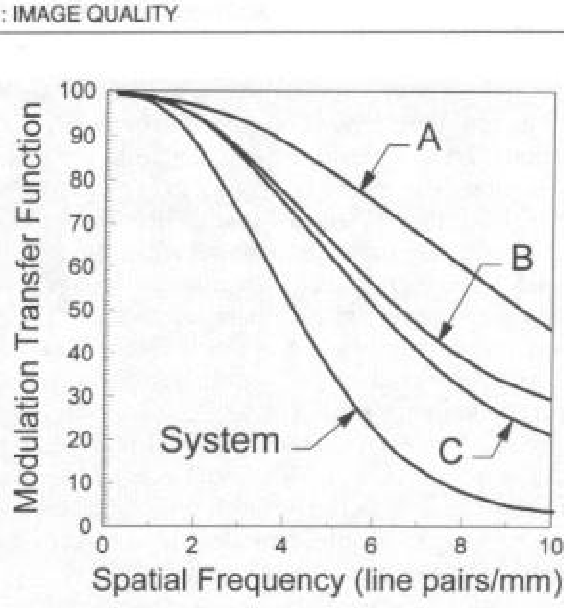

Figure 5.19. The MTF for several components in an imaging system. The total system MTF is also shown. The system MTF is determined by multiplying the other MTF curves (*A*, *B* and *C*) together at each frequency. The use of the MTF to describe resolution is very handy in this situation, because the weakest link in the imaging system (*C*) can readily be determined.

tifying which component is the weakest link allows designers to focus their attention where it will do the most good. It would be possible, for example, to measure the PSF for each of the components in a fluoroscopic system mentioned above. To calculate the resolution of the entire imaging chain, however, one would have to convolve all the PSFs of the subcomponents of the imaging chain. Because convolution is numerically intensive, this becomes difficult. The benefit of displaying the MTFs of all the components in the system is that the system MTF is simply the product of the subcomponent MTFs that make up the imaging system. Multiplication is simpler to perform than convolution, but more importantly, it is also easier to visualize in one's mind than the convolution of the individual PSFs. The cascade of the resolution through an imaging system is illustrated in Figure 5.19.

The Fourier Transform

The Fourier Transform is a mathematical operation developed by a French mathematician in the early 1800s. Although perhaps not essential to the everyday life of a radiologist, the Fourier transform is a routine part of life for the image scientist. The Fourier transform allows the conversion from the spatial domain to the spatial frequency domain. The inverse Fourier transform (FT^{-1}) can be used to convert back. Such manipulation is necessary because, if the Fourier transform is taken of the line spread function, one gets the MTF.[1] This approach is used in the rare

[1] Technically, the Fourier transform of the line spread function is the Optical Transfer Function, which is a complex function (i.e., it has imaginary and real components). The MTF is the modulus of the OTF, $MTF = \| OTF \|$.

instance in the radiology department where the MTF of an imaging system must be calculated precisely (as opposed to being estimated for routine quality control), for instance, when comparing screen-film systems from various vendors. The method of choice is to measure the LSF of the imaging system with a "slit camera," digitize the LSF (if not digital already), and then compute the Fourier transform of the LSF using a computer. The result is the MTF. The Fourier transform is also used in image processing. Because convolution can take even a fast computer a fair amount of time, in some instances it is much faster to use Fourier transforms to perform the operation. The Fourier transform is applied to both the image and the kernel it is to be convolved with, and then these are multiplied together. The inverse Fourier transform is then performed on the product. The result, now in the spatial domain, is mathematically identical with the convolution between the image and the kernel. This technique can be used to restore degraded images, or to apply a desirable appearance to an image.

Noise Frequency

A histogram of the noise distribution (see Fig. 5.6) does not clearly portray all the characteristics of the noise in an imaging system. Another property that noise has is its spatial frequency. The MTF is a graphical description of the resolution properties of an imaging chain, which describes the passage of a *signal* through the imaging chain. When *noise* is passed through the imaging system, its spatial frequency distribution is altered by the imaging system. When the Fourier transform is computed on a LSF signal, the MTF results. When the Fourier transform is computed on an input that has no contrast (and hence only noise, e.g. an image without an object or patient present), a description of the frequency distribution of the noise is obtained. It is common to square this frequency distribution, and the result is called the *noise power spectrum* or the *Wiener Spectrum*.

SAMPLING AND ALIASING

When dealing with frequencies, either temporal or spatial, it is important that the measurement of the frequency is accurate. For example, when calculating the MTF from the LSF, if the LSF is measured on film then it has to be digitized. The act of digitizing the LSF results in the true analog signal being diced into a discrete number of samples. If the number of samples or the spacing between the sample points is selected incorrectly, then it is possible for the calculation of the MTF to be in error. Specifically, if the sampling is too sparse, higher frequencies that may be present in the original analog LSF will *bleed into* the lower frequencies. This is called *aliasing*. A real-life example will illustrate aliasing. Most people have witnessed a strange phenomenon while watching a western on television. The spokes of a wheel appear as though they are rotating slowly backwards, even when they are obviously rotating forward at a much higher rotational frequency. This visual effect is the result of insufficient temporal sampling on the part of the imaging system. A television signal

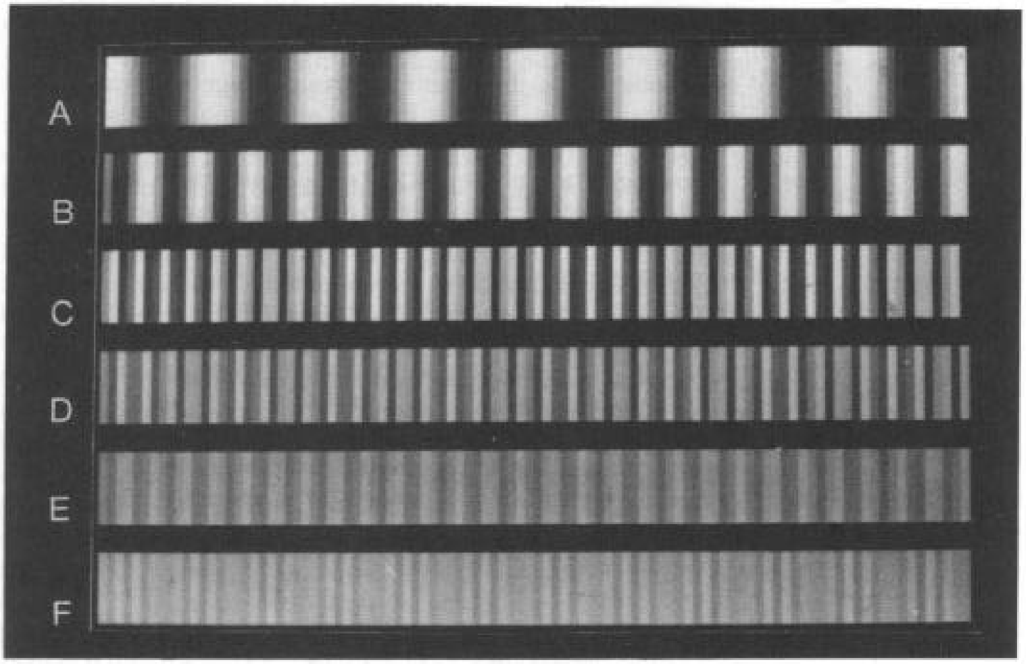

Figure 5.20. A visual example of aliasing. Aliasing occurs when the object that is being imaged has higher frequencies than the Nyquist frequency. The width of the six sinusoidal patterns is 1 cm, and there are 100 samples along this width, making the sampling distance 0.10 mm. The corresponding Nyquist frequency for a sampling of 0.10 mm is 5 line pairs per mm ($\Delta = 0.10$ mm, $fN = (1/2\Delta) = 5$ lp/mm). The six frequencies (before sampling) that are shown from top to bottom (*A to F*) correspond to 0.8, 1.6, 3.2, 6.4, 12.8, and 25.6 lp/mm. The top three strips adhere to the Nyquist limit ([0.8,1.6,3.2]<5.0), and the correct number of sinusoidal cycles can be counted (8, 16 and 32). The bottom three strips do not meet the Nyquist limit ([6.4,12.8,25.6]>5.0), and an incorrect number of cycles are displayed.

samples a new image every 1/30th of a second, a sampling *frequency* of 30 frames per second. If a wheel were moving faster than *half* of the sampling frequency (i.e., anything faster than 15 revolutions per second), aliasing can occur. *For aliasing not to occur, the sampling frequency has to be greater than twice the value of the highest frequency in the signal.* This "law of sampling" is called the *Nyquist criterion.* Aliasing can occur whenever there are frequencies in the signal that exceed the Nyquist limit. A visual example of aliasing is shown in Figure 5.20.

RECEIVER OPERATING CHARACTERISTIC CURVES

The direct measurement of parameters that influence image quality, like the SNR or the MTF, result in hard, objective data that describe the performance of the imaging system. However, they are not measures of diagnostic performance in the literal sense. To evaluate, for example, the ability of a class of images to reveal a certain pathology, more subjective measures are often used. This type of analysis involves asking if a person actually sees something abnormal in the image, and with what level of confidence, for a large set of images. Often radiologists are used as participants in a study that involves radiologic diagnosis; however, if a straightforward object (such as a sphere) is imaged, participants unfamiliar with radiologic imaging can be used. This type of image perception study is not as objective as the other measures of image quality discussed previously in this chapter, because there is a fair amount of subjective, human judgment involved in the evaluation. Although more subjective, receiver operating characteristic (ROC) studies provide a more direct measure of the performance capabilities of the imaging system or technique being evaluated, ROC analysis is particularly valuable for comparing different imaging procedures.

To perform an ROC analysis of an imaging system, a set of images must be compiled (or acquired) that demonstrates both normal and abnormal pathology. For example, two screen-film systems are to be compared for use in mammography. The test object is a small set of aluminum (Al) filings that are meant to simulate microcalcifications. The test object should be selected so that it is at the threshold of being visualized; blatantly obvious or impossibly faint test objects are not useful. A series

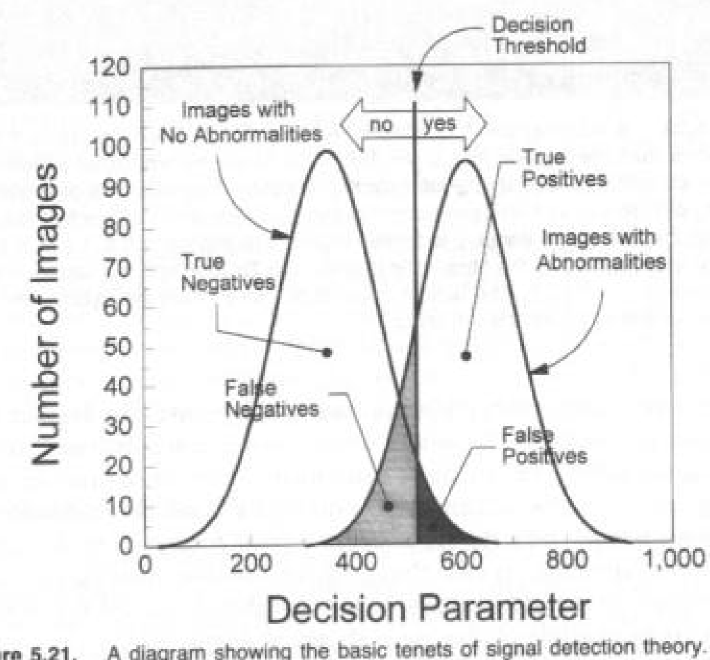

Figure 5.21. A diagram showing the basic tenets of signal detection theory. Images with no abnormalities are represented by the histogram of the left; images with abnormal features are represented by the histogram to the right. Because noise in these images gives breadth to each of the distributions, they overlap. The overlap implies that some of the images without an abnormality will appear as if they have one and that some images with abnormalities will appear as if they have none. This erroneous appearance is due to noise and its perception in the detection process. At some point, one must apply a decision threshold to state whether one observes an abnormality (this is actually a very subjective step). This decision threshold can be shifted left or right along the x axis. If it is shifted left, the number of false negatives will be reduced (the detection of actual abnormalities [sensitivity] will improve), but the number of false positives will increase (the determination of no abnormality being present when there is in fact no abnormality present [specificity] will be reduced.).

of 100 images is acquired using each screen-film system, system A and system B. In each series, the Al filings were present in the image about 50% of the time. It is essential to know whether the "pathology" was actually present in each image. In our example, the "truth" is known for each image because the presence or absence of the test object (the Al filings) was recorded as each image was acquired.

A participant in the study is presented each image sequentially, and asked to rank his or her confidence that the "microcalcifications" were there. It is typical to use a scale such as "definitely not present," "probably not present," "uncertain," "probably present," "definitely present," which is essentially the same as ranking one's confidence on a 1 to 5 scale. Each of these different levels of certainty represents a different *decision threshold*. The choice of the decision threshold determines the tradeoff between *sensitivity* and *specificity* (Fig. 5.21). In such an academic exercise as an ROC study, different decision thresholds are applied consciously. In the clinical interpretation of images, however, decision thresholds are routinely used by the radiologist, often unconsciously. For instance, in mammography, if 99% of biopsies of suspicious lesions called by a radiologist turn out to be benign (an example of poor specificity), the radiologist may wish to adjust his or her decision threshold to reduce the number of biopsies. Nevertheless, one should not adjust too far in the other direction, because this would result in missing too many actual cancers (a reduction in sensitivity), with potentially more serious consequences for the patient. Other factors act to influence the decision threshold, for example, the age of the woman, her familial history, and whether she had other symptoms.

After the ROC participant has ranked his or her confidence of the Al filings being present, a series of 2×2 decision matrices are generated, such as in Figure 5.22. At each decision threshold, sensitivity and speci-

Observer Says Lesion is There?	Lesion Actually There? YES	Lesion Actually There? NO
YES	TP	FP
NO	FN	TN

Figure 5.22. At a given decision threshold shown in Figure 5.21, the number of true positives (TP), false positives (FP), false negatives (FN), and true negatives (TN) can be determined. The 2×2 decision matrix shown here demonstrates the definitions of TP, FP, FN, and TN.

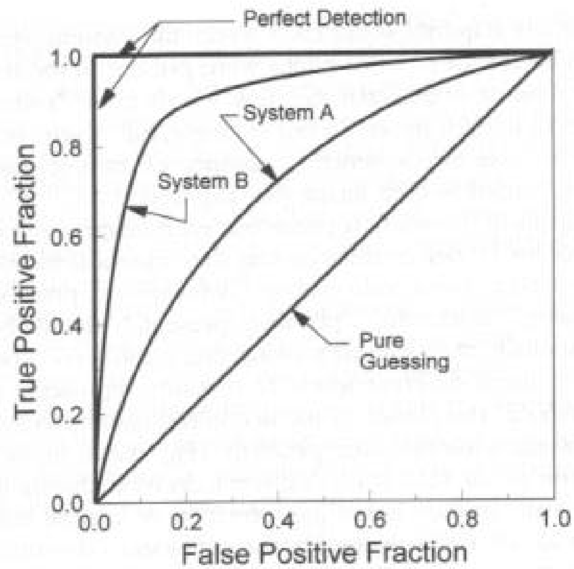

Figure 5.23. A Receiver Operating Characteristic (ROC) curve. A single decision threshold (Fig. 5.21) allows the assessment of TP, TN, FP, and FN as defined in Figure 5.22. From these values, the false-positive fraction is calculated as FP/(FP + TN), and the true-positive fraction is calculated as TP/(TP + FN). The true-positive fraction is the sensitivity, and the false-positive fraction is equal to one minus the specificity. An ROC curve is a plot of the sensitivity versus one minus the specificity. A single decision threshold would produce only a single point along the ROC curve. The entire curve is generated by sliding the decision threshold along the decision parameter axis (see Fig. 5.21). At each position of the decision threshold, a separate sensitivity and specificity can be calculated. An ROC curve is a plot of all those points linked together. Because an ideal imaging system has 100% sensitivity and 100% specificity, it is represented by a line running along the left and top borders of the ROC plot. Pure guessing would result in the diagonal line shown. Better imaging systems have curves that come closer to the upper left corner in the ROC curve. System B is therefore better than system A, as shown here.

ficity can be calculated. *Sensitivity* is the ability to detect a lesion when it is really there, whereas *specificity* is the ability to say that the lesion is absent when it is really not there. An ROC curve is a plot of the sensitivity versus (1-specificity), at each decision threshold. After an individual has evaluated both screen-film systems, these ROC curves are plotted on the same graph for comparison. Such an ROC curve is shown in Figure 5.23. If the participant closed his or her eyes for the test and randomly guessed whether the test object was present, a diagonal line representing "pure guessing" would result. In this figure, the results of an ideal system are shown, illustrating 100% sensitivity and specificity. This illustrates either an excellent screen-film system, or a test that was performed on objects that were blatantly obvious. Because the ROC test depends on the selection of the test object, they are used primarily for comparing two or more systems. ROC tests can also be used to compare the abilities of different observers, e.g., two or more radiologists, or a radiologist and family practitioner.

SUMMARY

In this chapter, the concepts important in image quality have been discussed generically. Of course, for each imaging modality, there are many nuances that are important but have not been mentioned here. By knowing what factors of a particular imaging technology influence the image quality, the radiologist can take steps to improve or customize the image quality depending on the requirements of the clinical examination.