Chapter 5
Image Quality

5.1 Introduction

Most of us can look at an image and immediately make some sort of judgement about its quality. We might notice that the image is blurry, or the colors are unnatural, or we can't read some text in the image because it is almost the same color as the background. These *subjective* assessments of image quality are made in relation to either the purposes of the viewer or, possibly as a secondary thought, the purposes intended by the image's creator. In visual art, at least, there may be a huge intellectual gap between those two purposes. In the case of medical images, however, the purposes of the image creator and viewer should be identical and well-defined. The purpose of a medical image is to provide information about *specific* aspects of body structure or function. The information in the image, in other words the *quality* of the image, must be sufficient to permit an accurate assessment of the structures or functions in question. If a doctor requests an X-ray for a suspected fracture then we want an image that will reveal not just the presence of a fracture but also its extent and severity. Such information may affect the way the patient's injury is managed. The radiologist reviewing the images will have very firm ideas about what constitutes a good quality image for the purpose of detecting and describing a bone fracture.

Since the purposes of medical images are objectively well-defined it is desirable to have quantitative methods of describing image quality. These methods are essential to the objective comparison of images and imaging systems, and the optimization of imaging technique and system design. Technique optimization is especially important when, as in X-ray imaging, there is an essential conflict between image quality and potential harm to the patient. Imaging cost is also an important consideration. Medical imaging systems are expensive to buy and operate so it is important to produce the best quality images in the shortest possible time.

In general we can say that the most important factors contributing to medical image quality are *Contrast*, *Spatial Resolution*, and *Noise*. *Ideally* we would like high contrast, high spatial resolution, and low noise, however, these are not independent factors – they affect each other in complex and confusing ways. If either contrast or spatial resolution is too low, or if noise is too high then, as illustrated in Fig. 5.1, the

R. Bourne, *Fundamentals of Digital Imaging in Medicine*,
DOI 10.1007/978-1-84882-087-6_5, © Springer-Verlag London Limited 2010

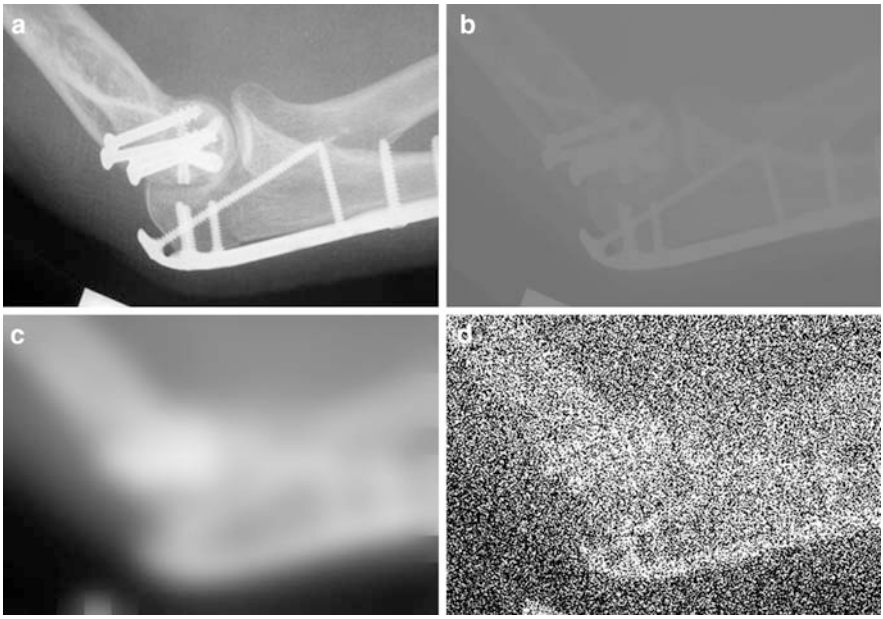

Fig. 5.1 A useful image must have adequate contrast and resolution, and a low noise level, as illustrated in the image (**a**). Image **b** has high spatial resolution and low noise, but is rendered useless by having almost zero contrast. Image **c** has low noise and high contrast, but extremely poor spatial resolution. In image **d** we see high spatial resolution but the very high noise level has destroyed the contrast information

image is of no value – no clinical information can be extracted. The actual levels of contrast and resolution that we need, and noise that we can tolerate, depends on the purpose of the image. They also depend on innate properties of the object imaged (mostly tissue), the properties of the imaging system (hardware), and the way the system is used (technique).

In this chapter we will first discuss contrast and noise separately and then, because noise strongly affects contrast, look at them together. We will discuss spatial resolution last because it is affected by both contrast and noise.

5.2 Contrast

All imaging techniques depend on *differential* emission of energy from the imaged object according to some physical property of the object. Without differential emission the measured signal would contain no information about the imaged object and the only variations would be informationless noise. *Contrast* is a measure of the magnitude of the measured signal differences between physically different regions

of the imaged object. When these measured signals are converted into an image 'contrast' describes the magnitude of intensity differences between different regions in the image. We can thus think of image contrast as being the product of *signal contrast* and *detector contrast*:

$$C_I = C_S \times C_D \tag{5.1}$$

The signal contrast (C_S) depends on the energy source and the physical properties of the imaged object – it describes the range of energies emitted by the object. The detector contrast (C_D) depends on the way the signal emitted by the imaged object is modified (e.g. with a scatter suppression grid), detected, and recorded. We need *both* signal and detector contrast to form image contrast.

In digital imaging it is very easy to enhance or reduce image intensity differences in order to make them more or less obvious to a human observer. This is the process of *contrast adjustment*. However, there is no way to enhance contrast if there is no difference in the measured signal or raw data. Also, a simple contrast adjustment that exaggerates (amplifies) recorded signal differences will also exaggerate any noise. The result may be no improvement in the ability to extract information from the image.

5.2.1 Simple Measures of Contrast

Figure 5.2 shows two version of a wrist MR image. Image a plainly has higher contrast than image b – the bones are much brighter *relative to the surrounding tissue background*, even though in image b the average brightness of the bones is greater than in image a. How could we measure these contrast differences *quantitatively*?

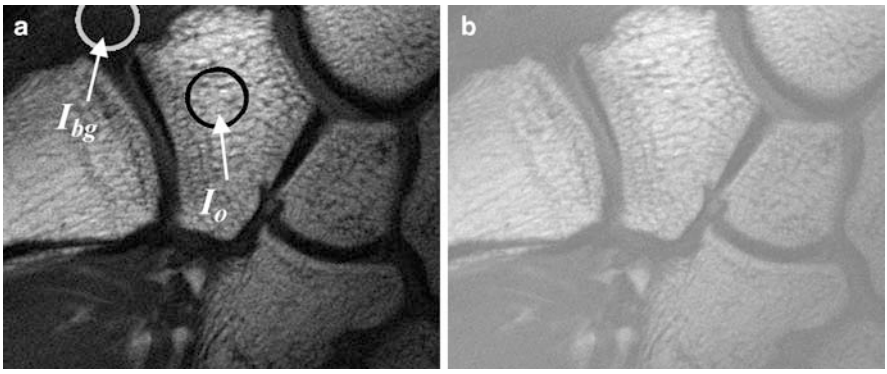

Fig. 5.2 Two versions of a wrist MR image. Image **a** plainly has higher contrast than image **b** – the bones are much brighter relative to the surrounding tissue background, even though in image **b** the average brightness of the bones is greater than in image **a**. Measures of contrast are based on the relative or absolute difference in average intensity of an object and its background

A simple way to do this would be to measure the average pixel intensity in a 'typical' bone region (inside the black circle, say) and in a typical soft tissue 'background' region (inside the white circle). The difference between these average intensities is 124 for image a and 63 for image b. This measurement suggests that the contrast in image a is about twice the contrast in image b. This figure seems a little out of kilter with our perception which would suggest *much* greater contrast difference between the two images. Another problem with this measurement is that we don't have a meaningful measure of the contrast in a single image. The intensity difference of 124 measured from image a is based on an 8 bit scale ranging from 0-255. If we had a 10 or 12 bit image we would get different measures for the intensity difference.

A more general measure of contrast describes the intensity difference relative to the background:

$$C = \frac{I_o - I_{bg}}{I_{bg}} \tag{5.2}$$

where I_o and I_{bg} are the average pixel intensities in the object and its background.

Using this measure we would get contrast values for bone versus background of 3.35 for image a and 0.44 for image b, suggesting that the contrast difference in image a is 7.6 times greater than in image b. This measure perhaps exaggerates the difference relative to our perception. Also, sometimes it is an arbitrary choice as to which region is the 'object' and which the background, yet the choice may make a significant difference to the contrast calculated by Eq. 5.2. We can avoid this problem by using the expression:

$$C = \frac{I_o - I_{bg}}{I_o + I_{bg}} \tag{5.3}$$

giving contrast values 0.63 for image a and 0.18 for image b, or a contrast ratio of 3.5 which is in good agreement with perception. Using this expression the choice of background and object affects only the sign of the calculated contrast. We will see soon that this expression is very similar to the definition of signal 'modulation'.

Another common perception-related way to measure contrast is the log of the intensity *ratio*:

$$C = log_{10} \frac{I_o}{I_{bg}} \tag{5.4}$$

For a film image this definition of contrast is identical to the *optical density* difference ($OD_o - OD_{bg}$). Human visual perception of intensity is non-linear (see Fig. 6.20) so this logarithmic scaling makes sense. For Fig. 5.2 it gives bone to soft tissue contrasts of 0.62 for image a and 0.16 for image b, or a 3.9 times difference between images a and b – a reasonably close match with our perception.

5.2.2 Contrast and Spatial Frequency

Human perception of contrast is dependent on spatial frequency. For a sinusoidal contrast pattern (Fig. 5.3) maximum contrast sensitivity occurs at a spatial frequency around three cycles per degree. Here a degree refers to an angle in the visual field. If we view an image at a distance of 400 mm, then one degree is 7 mm, so three cycles per degree = one cycle per 2.33 mm. Contrast sensitivity is about four times lower at frequencies of 0.5 and 8 cycles per degree, and 100 times lower at 40 cycles per degree.

5.2.3 Optimizing Contrast

The choice of imaging modality and imaging technique are fundamentally decisions about contrast because contrast is the image property that contains information. Spatial resolution describes the bottom limit of the scale of information potentially available, but as Fig. 5.1 clearly demonstrates, high spatial resolution is meaningless in the absence of contrast. If contrast in the measured signal cannot be obtained by virtue of the inherent physical properties of the imaged tissue and selection of an appropriate energy source, then it is often added artificially. A radiographer 'gives contrast' by injecting, or getting the patient to drink, a modality-specific contrast agent – an X-ray absorbing metal such as barium, a paramagnetic metal for MRI, or 'microbubbles' for ultrasound imaging.

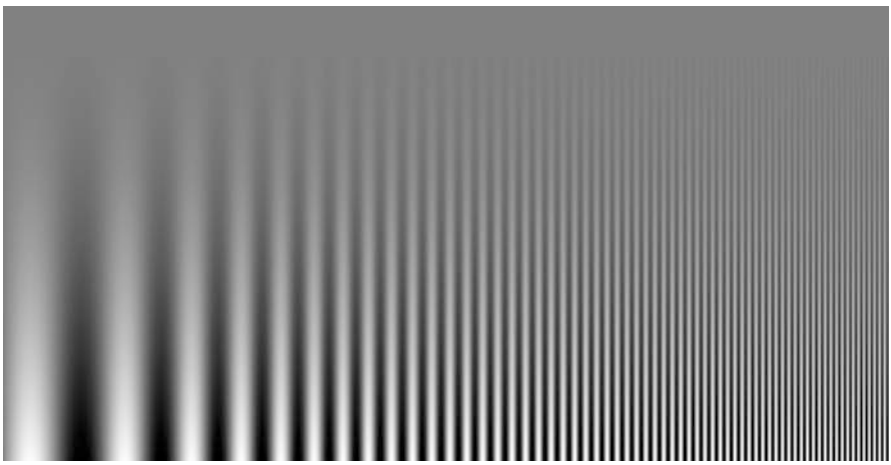

Fig. 5.3 Illustration of the spatial frequency dependence of contrast sensitivity. Contrast sensitivity is greatest at a spatial frequency of about 3 cycles per degree of visual field (≈1 cycle/3 mm at a viewing distance of 500 mm). Viewing this image at different distances will demonstrate the effect – the highest point of the bar pattern will move to the right as viewing distance decreases

5.3 Image Noise

No imaging method works without contrast and no imaging method is free of noise. If contrast is low and noise is high then the random intensity variations due to noise will make it difficult to visually detect the intensity changes due to contrast. Figure 5.4 illustrates how an increasing level of noise relative to contrast diminishes our ability to distinguish objects in an image. But before discussing the interrelation between noise and contrast we will first consider the origin and types of noise found in medical images.

5.3.1 What Is Noise?

The everyday answer to this question is *'Annoying sounds that make it difficult to hear what you want to hear'*. In terms of imaging a similarly broad definition would be *'Any intensity or color fluctuations that make it difficult to see what you want to see'*. The problem with these definitions is that we often don't *know* what it is we are supposed to be hearing or seeing – some or even all the information is hidden by the noise. What we want ideally is to be confident about the information we

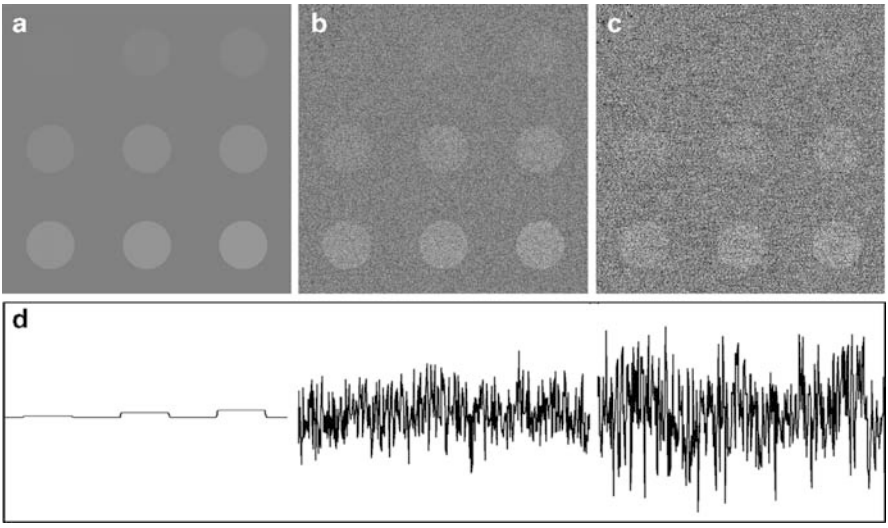

Fig. 5.4 The presence of noise reduces our ability to extract information from an image. The noise-free test pattern (**a**) contains nine circles of equal diameter with contrast ranging from 0.003 to 0.069 according to Eq. 5.4. In the presence of noise (images **b** and **c**) the visibility of the lower contrast objects is reduced or completely lost. Plot **d** shows the intensity profile through the top row of circles. Notice that objects are easier to distinguish in the image than in the intensity profile. This is because our perception computes a spatial average of intensities in the (2D) image but not in the 1D profile

receive. If you were on a very noisy phone connection you might ask your caller to repeat a sentence three times before you are sure you know all the words in the sentence. You may never hear one sentence fully due to the noise but so long as the caller repeats the same sentence each time you eventually know what all the words are. When we make an image we can effectively do the same thing to reduce the uncertainties due to noise – we either measure the signal for a longer time, or we repeat the measurement several times, which is effectively the same. If the noise is random its contribution to the total image signal will diminish over time or repeated measurements.

What happens if the noise is not random? This type of noise does not decrease when the measured signal is averaged over time. It can appear due to leakage of a spurious signal into the system (from the mains power supply for example), or some systematic 'error' in the method of formation of the image. The latter type of signal is generally referred to as an *artifact*.

Focusing on imaging we can, at least conceptually, separate the energy we measure, the image (I), into 'true' signal variations that are characteristic of the imaged object (S), and 'spurious' or 'random' variations that are *not* characteristic of the imaged object – in other words noise (N) (and possibly artifacts but we will ignore these for now):

$$I = S + N \tag{5.5}$$

While we would generally think of the 'signal', S, as always being positive, while the noise, N, could be expected to make either positive, zero, or negative contributions to a measurement. The 'erroneous' measurement value it causes may be either higher or lower than the 'true' signal measurement expected. In many cases we can say that the average (mean) measurement error due to noise is zero. This is 'zero-mean' noise.

In general we have no way of separating 'true signal' from 'random noise'. Nevertheless, the concept of separable signal and noise is very useful in signal and image processing. This does, however, cause some potentially confusing terminology. When we talk about the intensity of the energy we measure or 'detect' this is usually referred to as 'the signal' – *'I'm not getting a signal!'* or *'The signal is strong enough, but it's very noisy'*. A few seconds later somebody will ask, *'What's the signal-to-noise ratio?'* Suddenly the meaning of 'signal' has changed. In discussions that include noise, 'signal' usually means the notional 'true' signal, free of noise, coming from the object of interest. If the discussion does not explicitly mention noise then 'signal' probably means the intensity of detected energy, including noise. Most of the time the context is sufficient to work out which usage is applicable.

5.3.2 Quantum Mottle

Going back to our general definition of image noise as 'Any intensity or color fluctuations that make it difficult to see what you want to see' we should include in this

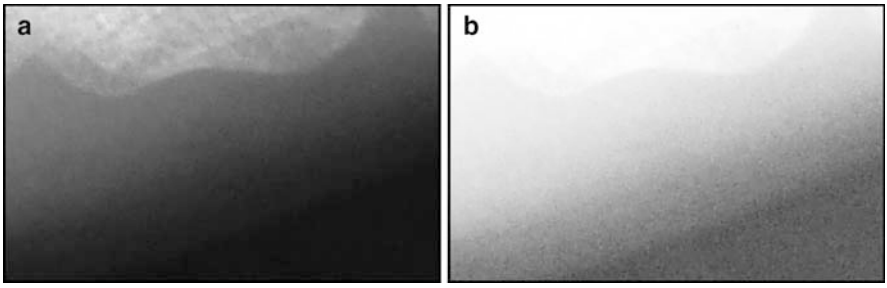

Fig. 5.5 The random 'grainy' pattern in the background of a radiograph (**a**) is called *Quantum mottle*. Here it is made more obvious in image (**b**) by contrast enhancement of the darker pixels. Quantum mottle is the natural result of statistical variations in the number of X-ray photons incident on a particular part of the detector in a finite measurement time interval

random fluctuations in the flow and spatial distribution of energy from the energy source itself. In the majority of medical imaging methods the energy source directly (or, in the case of PET, indirectly) emits electromagnetic radiation, or *photons*. The particulate nature of electromagnetic radiation has a profound effect on the nature of noise in these medical images. The probability of a photon being detected in any region of the image sensor during any specific interval of time is *constant*. We can see this effect if we look closely at the background of a radiograph (Fig. 5.5) where the variation in measured intensity is due entirely to the (locally) random distribution of photons emitted by the X-ray source. This random 'texture' is called *Quantum mottle*.

Quantum mottle is the result of a *Poisson* process, meaning that if we measure the standard deviation (σ) in the signal for all identical sensor regions (pixels in the case of a digital sensor) it will be equal to the square root of the average signal (μ). It is helpful to describe this variation in relative terms:

$$\textit{Relative variation} = \frac{\sqrt{\mu}}{\mu} = \frac{1}{\sqrt{\mu}} \tag{5.6}$$

The significance of this expression is that the more accumulated signal that we get, the smaller is the relative variation. In other words, quantum mottle *decreases*, relative to total intensity, with increasing measurement duration or signal energy flux. If we measure for four times as long, or take the average of four measurements, then the amount of quantum mottle noise will decrease by one half. Poisson statistics is applicable to a large number of physical processes occurring in medical imaging.

5.3.3 Other Noises

Quantum mottle is an example of 'measurement noise' that is an inevitable consequence of the use of electromagnetic radiation as the imaging energy source. Even

when we have reliable statistics in the detected energy there are other ways noise can enter the recorded signal and thus appear in an image. The most obvious of these noise sources is thermal noise in the sample and the electronic hardware associated with signal detection, amplification, and recording. Also, spurious electromagnetic radiation from external sources can leak into the analog electronics. The digital components of an imaging system are much less susceptible to external noise but can nevertheless give rise to random variations in reconstructed images due to lack of calculation precision.

5.3.3.1 Gaussian Noise

Random noise that enters the imaging system from external sources typically has a Gaussian, or normal, distribution. The magnitude of the 'error', or uncertainty, in the recorded signal due to this type of noise is not dependent on the total signal, as is the case for quantum mottle. In an image Gaussian noise affects dark and light areas to the same degree (Fig. 5.6a).

The intensity uncertainty due to Gaussian noise can be described by a *Probability Distribution Function* (PDF) shown as the familiar 'bell curve' in Fig. 5.6b. In imaging terms this PDF tells us there is a high probability that, for any particular pixel, the error due to noise will be small. Large errors have very low probability. The average of all errors is zero because the noise can either increase or decrease the recorded signal. The standard deviation (σ) of the PDF gives an indication of the average error *magnitudes*. For a Gaussian distribution $\frac{2}{3}$ of all pixels will have an error less than (σ). It is important to remember, however, that although we may be able to summarize the error statistics there is no way to know how big the error is for any particular individual measurement (pixel). Thus there is no way to completely remove the noise.

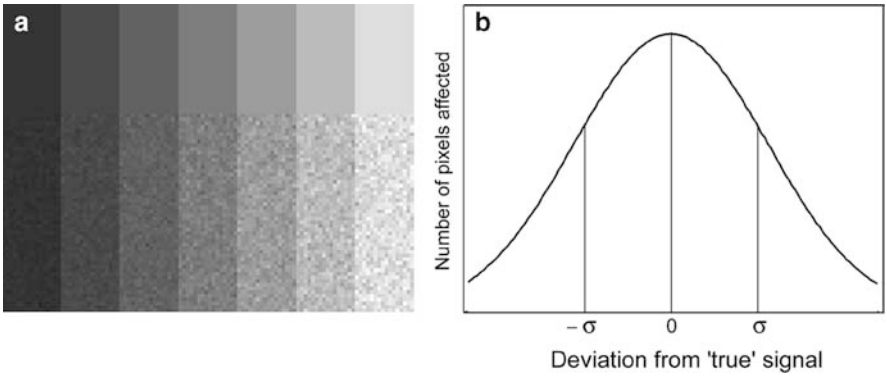

Fig. 5.6 Gaussian noise affects light and dark areas of an image equally. Image **a** shows a test pattern without and with Gaussian noise. The *Probability Distribution Function* (PDF) (**b**) describes the distribution of the noise errors. The standard deviation (σ) gives a summary of the range of error magnitudes. For an image with a Gaussian noise distribution, $\frac{2}{3}$ of all pixels will have an error less than (σ). A decrease in Gaussian noise is equivalent to a decrease in σ

Gaussian noise is sometimes called 'additive noise' as its contribution to an image, I, can be modeled as a simple addition of noise having a Gaussian PDF, N_G, to the noiseless signal, S:

$$I = S + N_G \tag{5.7}$$

5.3.3.2 Speckle

Speckle is the dominant form of noise in medical ultrasound imaging (Fig. 3.14) because many microscopic tissue components, smaller than the spatial resolution of the technique, act to reflect sound waves incoherently. When waves of a single wavelength are reflected from a rough surface each point on the surface effectively acts a source of spherical waves. The roughness of the surface means that these numerous spherical waves will have many different phases which will cause the waves to interact *constructively* at some points in space and *destructively* at other points. If the surface is rough enough to cause phase differences greater than 2π (360°) then the resultant intensity pattern will be random. The maximum possible intensity would be the sum of the amplitudes of a large number of large amplitude spherical waves, whereas the minimum intensity, when all waves interact destructively, would be zero. High intensity constructive interactions are much more probable in regions where numerous high amplitude spherical waves are present – in other words, in regions of high average signal intensity. In regions of low average signal intensity constructive interactions will produce only medium or low amplitudes. The overall result is that the random (noise) signal will have an average amplitude that increases with overall signal intensity. In an image this appears as many bright specks in the lighter regions. This is *Speckle* noise.

Speckle is not an *entirely* bad thing in ultrasound. The speckle pattern is not always completely random – in which case it contains information about the tissue. Some pathologies give rise to a characteristic speckle signal. In contrast to Gaussian 'additive' noise, speckle noise is called 'multiplicative noise'. It's contribution to an image, I, can be modeled as multiplication of the noiseless signal, S, by random numbers having a zero-mean Gaussian PDF, N_G:

$$I = S + (S \times N_G) \tag{5.8}$$

Note that in this model we are still adding the noise to the 'true' signal, but this time the magnitude of the noise is proportional to the signal intensity.

5.3.3.3 Salt and Pepper

Most image noise causes relatively small random deviations from the expected 'true' image intensity. Extreme intensity deviations may also occur resulting in image pixels being incorrectly assigned the maximum or minimum possible values. This 'Salt

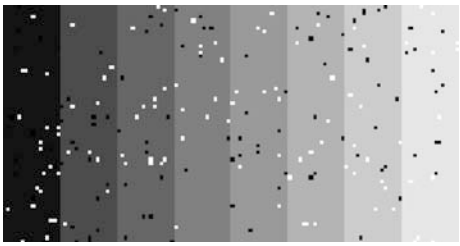

Fig. 5.7 'Salt and Pepper', or *Impulse*, noise appears as randomly distributed *white* and *black* pixels in a digital image

and Pepper', or *Impulse*, noise appears as randomly distributed white and black pixels in a digital image (Fig. 5.7). It can arise from defects in single elements of a semiconductor sensor, sometimes called 'stuck pixels', defects in a semiconductor memory or storage device, noise affecting data transmission, or even image reconstruction and processing errors.

5.3.3.4 Artifacts

Our broad definition of image noise as *'Any intensity or color fluctuations that make it difficult to see what you want to see'* would include systematic intensity errors that occur due to specific properties of the imaging method, or the interaction of these properties with those of the imaged object. Systematic errors of this kind are referred to as *artifacts*. We previously mentioned the 'chemical shift artifact' in MRI that causes the image of fat to be displaced relative to adjacent tissue. Because the origin of artifacts is highly specific to the imaging modality, rather than a general characteristic of the flow and differential transmission of imaging energy, the solutions are also very modality specific and we will not investigate them further in this text.

5.3.4 Signal to Noise Ratio

As we saw in Fig. 5.4 noise can have a significant effect on our ability to see objects in images. A common way to quantify the level of noise in an image (or any measured signal) is to estimate the *Signal-to-Noise Ratio* (SNR):

$$SNR = \frac{S}{N} \tag{5.9}$$

Fig. 5.8 A simple estimate of *Signal to Noise Ratio* (SNR) in an X-ray image. The *Signal* is measured as the average intensity in a selected region of the anatomy, in this case the base of the humerus, and the *Noise* in a region of the image where there is no signal from anatomy. We can then calculate (using Eq. 5.10) that, for the base of the humerus, the SNR is $0.66 \times \frac{152.9}{1.23} = 77.6$. Obviously, this method of measuring SNR is highly dependent on where in the image the signal is measured and it takes no account of the importance of contrast (the 'Results' windows are the output of the ImageJ Menu: Analyze > Measure command applied after drawing the ROI)

In this expression we would not want to use the mean value of the noise as, at least for zero-mean Gaussian noise, this would give an inappropriate infinite SNR value. Instead it is conventional to use the standard deviation of a tissue-free region in the background(σ_{BG}) as an estimate of the noise. The 'signal' can be taken as the mean intensity (μ_S) in the anatomy of interest (Fig. 5.8):

$$SNR = 0.66 \times \frac{\mu_S}{\sigma_{BG}} \tag{5.10}$$

A 'correction factor' of 0.66 is applied because in the background the 'negative' part of the zero-mean noise appears black and the measured standard deviation is less than the true value expected for a Gaussian PDF.

The 'difference method' is useful for measuring SNR of an imaging system rather than the SNR of a single image. In this method two images (I_1 and I_2) of a test object are acquired one after another without changing any imaging parameters. The 'signal' is then taken as the mean intensity (μ_S) in a ROI placed inside the image of the test object. The 'noise' is measured from the standard deviation (σ_D) in the same ROI in the *difference* image ($D = I_1 - I_2$). Notice that the difference image can have negative pixel intensities. The SNR is then calculated using the expression:

$$SNR = \sqrt{2} \times \frac{\mu_S}{\sigma_D} \tag{5.11}$$

In this case the correction factor compensates for the exaggeration of noise in the difference image.

5.4 Contrast + Noise

The problem with image SNR estimates is that they reveal nothing about the effect of noise on our ability to see objects in an image because visibility depends on contrast – the difference between signals. A highly overexposed radiograph might have a very high SNR and yet contain no useful information about the imaged object. A more useful estimate of the effect of noise on image information is the *Contrast-to-Noise Ratio* (CNR):

$$CNR = \frac{(\mu_A - \mu_B)}{\sigma_{BG}} \tag{5.12}$$

In this expression noise is measured as the standard deviation of the background and the contrast measure is simply the intensity difference between an object and its background.

Figure 5.9 shows a series of nine circles of increasing contrast relative to a mid-gray background. The contrast values range from 0.004 to 0.086 (expressed here as a fraction of the possible intensity range to eliminate the significance of bit depth), while the standard deviation of the noise is 0.173. This gives us CNR values ranging from 0.02 to 0.50. We can perhaps just barely see the circle at the left end of the middle row (CNR = 0.35).

The ease of visual detection of an object in an image, sometimes called 'conspicuity', depends not just on the CNR but also on the size of the object. Figure 5.10a shows a series of various size circles of constant CNR relative to the background. The largest circle is easy to see but the smallest, at the top left is effectively invisible. Because our perception performs a local 2D averaging of intensities larger objects are generally easier to see than smaller ones. The visibility or conspicuity of an object is roughly proportional to its area. However, for the same area, more compact objects (such as circles) will be more visible than less compact objects (such as stars) as demonstrated in Fig. 5.10b.

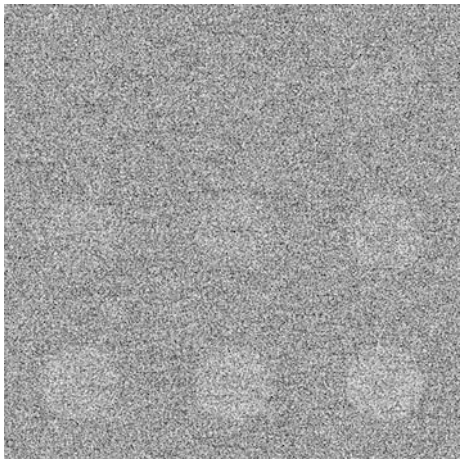

Fig. 5.9 Illustration of the effect of noise and contrast on object visibility. The image contains nine circles of equal diameter with contrast values ranging from 0.004 (*top left*) to 0.086 (*bottom right*), according to Eq. 5.3. The standard deviation of the noise is 17% of the possible intensity range. This gives CNR values ranging from 0.02 to 0.50. We can perhaps just see the circle at the left end of the middle row (CNR = 0.35)

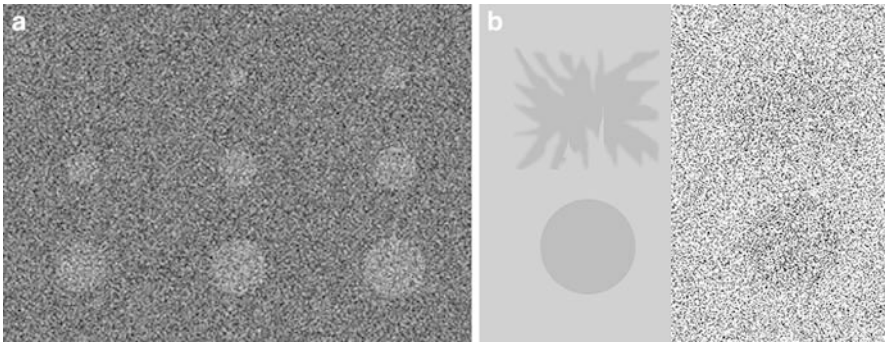

Fig. 5.10 Illustration of the effect of object size and shape on visibility in the presence of noise. The test image **a** contains nine circles of equal contrast but a range of diameters. Because our perception performs a local 2D averaging of intensities, larger objects are generally easier to see than smaller ones. However, for objects having the same area (e.g. the circle and the 'squashed butterfly' in image **b**), a more compact object (such as a circle) will be more visible than a less compact object

It should now be clear that contrast, noise levels, and object size and shape all affect our ability to extract visual information from an image. There are several ways of measuring SNR and CNR so we need to be careful to specify the method used when communicating such measurements.

5.5 Spatial Resolution

In an ideal situation 'spatial resolution' describes the size of the smallest objects that can be separately discriminated by an imaging system. However, as we have just seen, the effects of contrast and noise are very significant, so a proper description of spatial resolution needs to incorporate these factors.

5.5.1 Line Pairs

In plain X-ray imaging it is common to test and measure the spatial resolution of a system in terms of *Line pairs* per millimeter (or centimeter). The line pairs test typically uses a test object which comprises a series of parallel lines of very high inherent contrast (Fig. 5.11a). For a plain X-ray imaging the test object is a thin metal plate in which a series of slots of progressively decreasing width and spacing are cut. One 'line pair' in an image is then one white line immediately adjacent to one black line.

When a line pairs test object (phantom) is imaged in an imperfect system the acquired image is blurred to some degree (Fig. 5.11b). A crude measure of the spatial resolution of the imaging system is the maximum spatial density of line pairs that

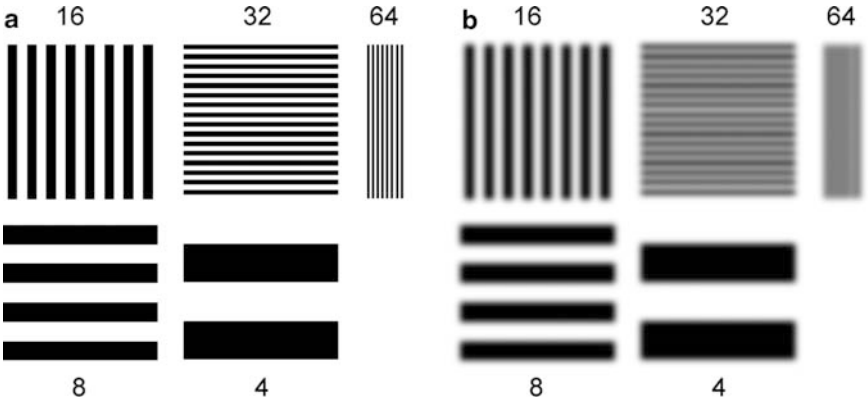

Fig. 5.11 A line pairs test pattern (**a**) as typically used to measure the spatial resolution of a projection imaging device. The numbers specify the number of line pairs per unit distance. Note the similarity to the black and white stripe patterns of Figs 4.10 and 4.11. Each line pairs pattern is described by a primary sinusoidal spatial frequency identical to the line pairs frequency. Higher spatial frequencies are also present but at lower amplitude than the primary. Image **b** shows a possible image obtained by imaging the test object in an imperfect system. Note that the effect of imperfect spatial resolution is a loss of contrast at higher spatial frequency

can just be distinguished. This will correspond, very roughly, to the smallest width *linear* object that would be visible in the image. Spatial averaging by our perception makes lines more visible that points of the same dimension. Point objects of the same diameter as the minimum visible line width will be *invisible*. Notice in Fig. 5.11b that blurring causes a progressive loss of contrast as the line pairs density increases. Even at line pair densities above the limit of visual resolution the light part of the pattern is no longer white and the dark part is not black – decreasing spatial resolution is due to a *loss of contrast*.

It should be clear from Fig. 5.11b that visual inspection of a line pairs test image is not an ideal measure of spatial resolution as there is no clear point at which separate lines become indistinguishable, and there is significant loss of contrast for objects much larger than the visibility limit. The effect of the blurring is a general loss of edge and line detail, and a progressively more severe loss of contrast with increasing spatial frequency. The *inherent* very high contrast of the test object is not fully reproduced by the imaging system – especially at high spatial frequency.

5.5.2 The Modulation Transfer Function

Rather than attempt to define spatial resolution with a single ‘minimum object size’ parameter it is better to fully describe the way the inherent contrast of an imaged object is lost in the imaging system as spatial frequency increases. This is commonly done with the *Modulation Transfer Function* or MTF. MTF descriptions of imaging system performance are used widely – in microscopy, photography, and astronomy,

to name a few. Applied to a medical imaging system, the definition of modulation (M) is very similar to the contrast definition used in Eq. 5.3:

$$M = \frac{I_{max} - I_{min}}{I_{max} + I_{min}} \tag{5.13}$$

where I_{max} and I_{min} are the maximum and minimum signal intensities (Fig. 5.12).

For a sinusoidal intensity change, as shown in Fig. 5.12, the modulation described by Eq. 5.13 is identical to the amplitude of the sinusoid (I_{amp}) divided by the background (I_{bg}) on which it is imposed. Notice that, by this definition, for image b the modulation M_B is greater than M_A for image a because, although the amplitudes are identical, the background intensity is lower in image b.

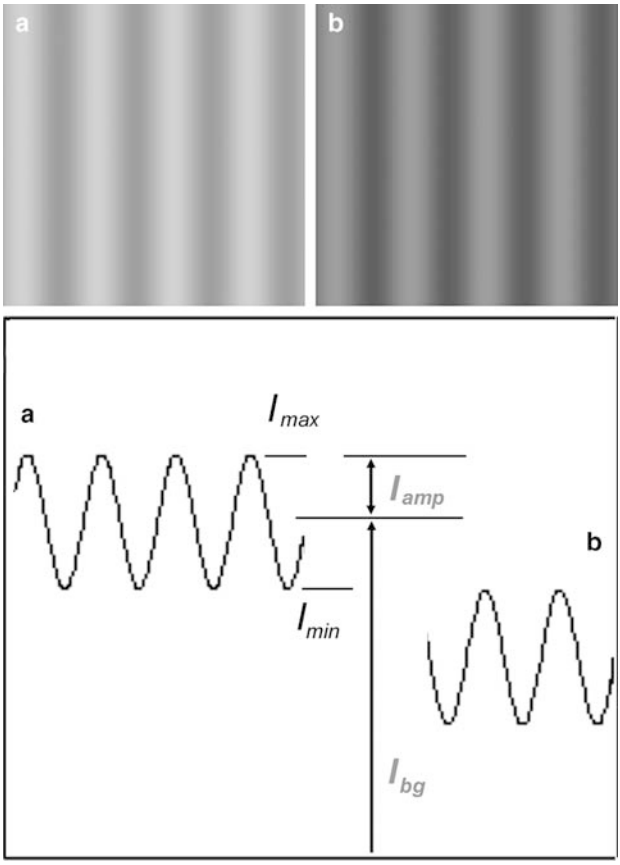

Fig. 5.12 The modulation (M) of image intensity is defined as $M = (I_{max} - I_{min})/(I_{max} + I_{min})$. For a sinusoidal intensity profile this definition is identical to $M = I_{amp}/I_{bg}$ where I_{amp} is the amplitude of the sinusoid and I_{bg} is the background intensity. Notice that for image **b** the modulation M_B is greater than M_A for image **a** because, although the amplitudes are identical, the background intensity is lower in image **b**

The MTF for an imaging system, or a component of an imaging system, is defined as:

$$MTF(f) = \frac{M_{out}(f)}{M_{in}(f)} \tag{5.14}$$

where $M_{in}(f)$ is the 'input' signal modulation and $M_{out}(f)$ is the output signal modulation. The symbol (f) is included to emphasize that we are interested in the way the input and output signal modulation changes with spatial frequency. When describing the MTF of a whole imaging system, M_{in} would be the inherent contrast in the imaged object, and M_{out} would be the image contrast.

For our notional imaging system that produced the blurred image shown in Fig. 5.11b we can estimate $MTF(f)$ by examination of the intensity profiles of the original and blurred line pairs patterns (Fig. 5.13). The intensity $(M_{in}(f))$ of the original pattern always varies between a maximum (I_{max}) of 255 and a minimum (I_{min}) of zero (this is an 8 bit image), so, according to Eq. 5.13, $M_{in}(f) = 1$ at *all* the line pair densities in the test pattern. $M_{out}(f)$, however, decreases with increasing spatial frequency. The five values of $M_{out}(f)$ at the spatial frequencies 4, 8, 16, 32, and 64 are approximately 1, 1, 0.8, 0.1, and 0 (we omit the space unit here as it depends on the final print size of the image). Notice that at line pair density 8 the intensity profile is quite rounded but I_{max} and I_{min} are still 255 and 0 respectively, so $M_{out}(8) = 1$. MTF is normally plotted, as shown in Fig. 5.14, against a linear or logarithmic spatial frequency scale.

In the above discussion we equated line pairs density with spatial frequency. This is a convenient approximation because, as we saw in Figs. 4.10 and 4.11, a line pairs pattern is described by a primary sinusoidal spatial frequency and its odd harmonics. The amplitudes of the harmonics are significantly lower than the amplitude of the primary frequency. For the estimation of MTF we need only consider the primary frequency.

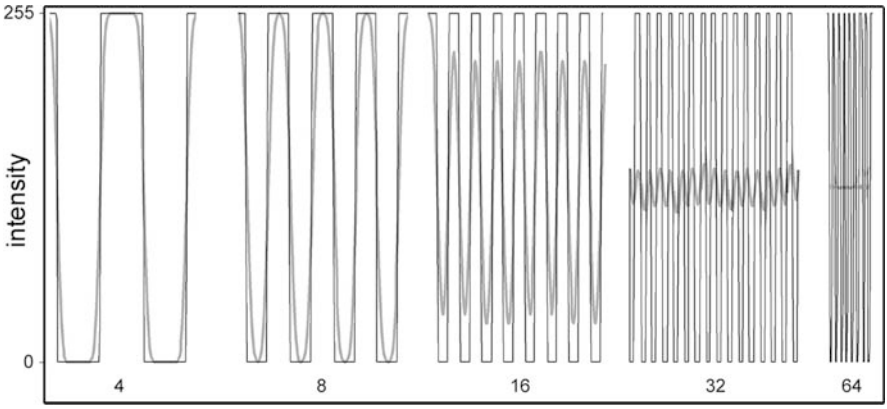

Fig. 5.13 Intensity profiles of the original (*fine black line*) and blurred (*bold gray line*) line pairs patterns from Fig. 5.11. In this example $MTF(f)$ simplifies to the ratio of amplitudes because $I_{max} + I_{min}$ is the same for both M_{in} and M_{out}

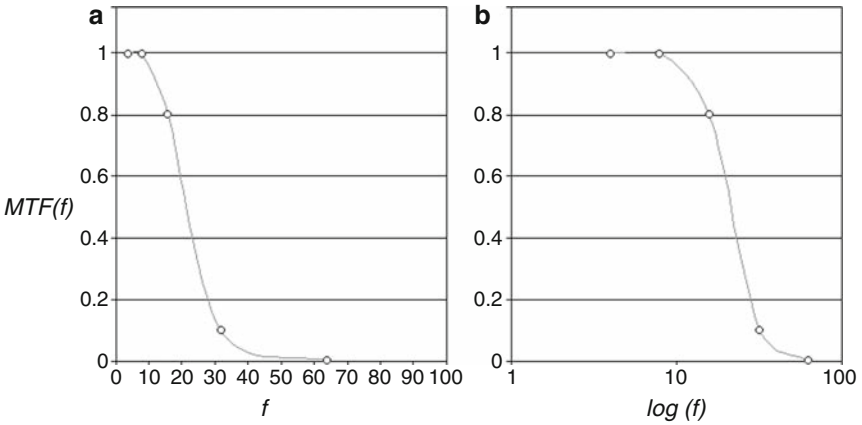

Fig. 5.14 Linear (**a**) and logarithmic (**b**) MTF curves that describe the blurring and loss of contrast in Fig. 5.11b

Recall from Chapter 4 (Fig. 4.27) that the roughly rectangular shape of an X-ray tube's focal spot resulted in slightly different blurring in the x and y directions. This corresponds to different $MTF(f)$ curves for the x and y directions. An $MTF(f)$ curve can be used to describe the spatial resolution performance of a whole imaging system, or specific components of an imaging system. Since all MTFs have a maximum value of 1.0 and a minimum of zero, the system MTF is simply the product of the MTFs of its individual components. Analysis of the system components separately enables the identification of the most likely targets for design improvements.

5.5.3 The Edge, Line, and Point Spread Functions

Instead of using a line pairs phantom to test an imaging system we could use an object that has an inherent contrast that varies sinusoidally in space. For an X-ray imaging system this might be a metal plate machined to a sinusoidal profile. Such a device would be extremely tricky to make, and not particularly useful, as it could only be used to measure MTF at a single spatial frequency. In fact we don't need a line pairs phantom either. All that is required is a high contrast object with a very sharp edge.

We saw in Chapter 4 that a sharp edge or narrow line feature in an image are both described by a large range of spatial frequencies in the Fourier spectrum of the image. In the limiting case of a line of width one pixel the amplitudes of all spatial frequencies are the same. In a blurred image the amplitudes of high spatial frequencies are small or zero. If we examine the Fourier spectrum of the image of an object known to have abrupt contrast changes, such as sharp edges or narrow lines, then we should be able to determine the degree to which the amplitudes of higher spatial frequencies are attenuated by the imaging system – in other words, $MTF(f)$.

$MFT(f)$ of an imaging system can be measured in any given direction by imaging the straight edge of a suitable test object, or a 'line' object. For an X-ray system we could use the edge of a metal sheet or a thin straight wire. The orientation of the edge or line should be perpendicular to the MTF direction of interest (remember that straight image edges produce a linear Fourier spectrum feature perpendicular to the edge). The imaging system will produce an image of the edge or line which is blurred to an extent specific to the particular imaging parameters and equipment used (Fig. 5.15a, b). A plot of the intensity profile perpendicular to the edge or line shows the spatial blurring in the direction of interest – the *Edge Spread Function* (ESF) or *Line Spread Function* (LSF) (Fig. 5.15c, d). Notice that the profile of the ESF is effectively half of the symmetrical LSF.

$MTF(f)$ can be obtained by plotting the 1D Fourier spectrum of the LSF, and scaling to a maximum of 1.0 (Fig. 5.15e). In this case we use the raw frequency domain amplitudes, not their logarithms, and the zero-frequency (DC) term is first discarded because we are only interested in the amplitudes of non-zero spatial frequencies.

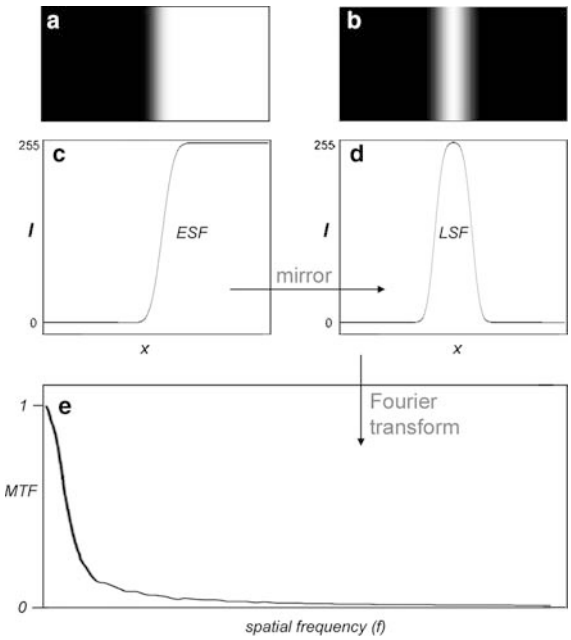

Fig. 5.15 The edge and line spread functions can be used to determine a direction-specific MTF. (**a**) Image of a sharp high contrast straight edge. (**b**) Image of a narrow line. (**c**) The profile of the intensity in image **a** describes the *Edge Spread Function* (ESF). (**d**) The profile of the intensity in image **b** describes the *Line Spread Function* (LSF). (**e**) The MTF is the 1D Fourier spectrum of the line spread function scaled to a maximum of 1. The ESF profile needs to be 'mirrored' to produce an equivalent LSF prior to Fourier transformation

An ESF needs to be converted to an LSF before Fourier transformation to prevent formation of an edge artifact. (The 1D Fourier transform treats its input data as if it repeated infinitely in one dimension – just as the 2D Fourier transform treats an image as if it were tiled infinitely in 2D space. Notice that the right side of the ESF is 255 and the left zero so there would be a large step in the repeated profile.)

We mentioned the *Point Spread Function* PSF several times previously as the essential spatial domain description of the 2D blurring tendency of an imaging system, or a specific component of an imaging system such as the focal spot of an X-ray tube (Fig. 4.27). Obtaining a 2D MTF is simply a matter of Fourier transforming the PSF and scaling the raw 2D Fourier spectrum (minus the DC term) to a maximum value of 1.0. The value of the 2D MTF derived from the PSF is that it describes the spatial resolution of an imaging system in all spatial directions – there is no need to perform a separate measurement for each direction of interest as is required when using the line or edge spread functions. However, obtaining an image that represents the PSF is not easy for imaging systems that have an inherently weak signal such as MRI scanners. The test object used to acquire a PSF image needs to have a 'point' feature smaller than the spatial resolution of the imaging system. For plain X-ray imaging this can be a pinhole in a thin metal sheet. For CT a wire aligned with the axis of the scanner will give a PSF in the axial plane. Similarly for PET and SPECT a small bead or wire of radioactive material is used to measure either a 3D or 2D PSF respectively.

5.6 Contrast + Noise + Resolution

From the preceding discussion it should be clear that our ability to extract information from an image is primarily dependent on the image contrast, while the spatial scale of the information available is dependent on the spatial resolution of the image. Noise reduces our ability to detect subtle contrast differences. Noise has a much less significant effect on spatial resolution due to the tendency of our perception to perform a 2D spatial averaging of intensities. Figure 5.16 provides an illustration of the combined effects of contrast and noise on spatial resolution. This image gives an artificially strong impression of spatial resolution because our perception tends to extend the linear contrast pattern into regions that, viewed in isolation would appear to contain only noise.

5.7 Summary

- The primary determinants of medical image quality are *Contrast, Spatial resolution*, and *Noise*. An ideal image has high contrast, high spatial resolution, and low noise.

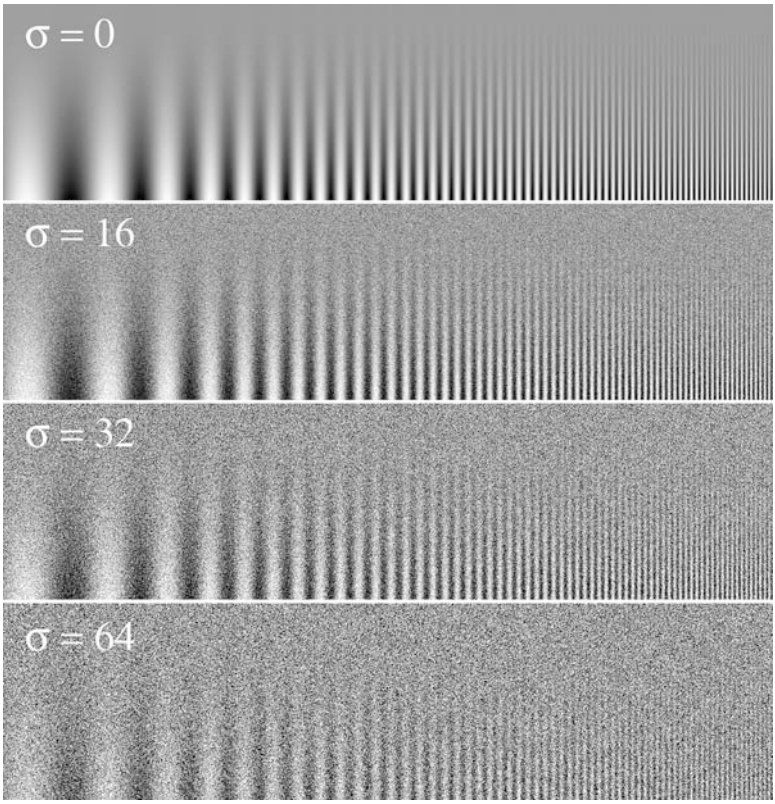

Fig. 5.16 Illustration of the effect of noise on contrast sensitivity at different spatial frequencies. The noise PDF in this test image is Gaussian with σ values expressed as 8-bit intensities (min = 0, max = 255). Noise decreases contrast sensitivity at all spatial frequencies but especially high frequencies

- Image contrast can be quantified as the absolute or relative difference between the average pixel intensity inside an object of interest and the average pixel intensity in the object's background.
- Most images are affected by noise – random signal variations that are superimposed on the 'true' signal of an imaged object. Noise decreases the contrast sensitivity of human perception. Low contrast objects are difficult to see in a noisy image.
- Noise can be characterized by a *probability distribution function* (PDF) – a summary of the statistical error due to noise in a recorded signal. Most image noise can be described by a *Gaussian* PDF (bell curve) with a characteristic standard deviation (σ). A Gaussian PDF means that 66% of all

image pixels have an intensity error due to noise that is less than σ. A more noisy image will have a larger value of σ.

- The energy detected by an imaging system can be thought of as being either *signal* (contains information about the imaged object), or noise (contains no information and obscures the signal). The *Signal to Noise Ratio* (SNR) is a measure of the quality, or *potential* information content, of the image data. If the noise energy is random then SNR can be increased by increasing the total amount of energy measured – either by applying a more intense flow of energy, or by measurement over a longer time period.
- The *actual* information content of a medical image usually depends on *contrast* – the differences in image intensity that reveal structural or functional properties of the imaged tissue. The *Contrast to Noise Ratio* (CNR) is often a more useful measure of image quality than SNR.
- The spatial resolution of an image is dependent on the construction and geometry of the imaging system, and the interaction of the imaging energy with the imaged object. All imaging systems blur the signal from the imaged object, and the object itself may scatter some of the imaging energy. The amount of blur is characterized by the *Point Spread Function* (PSF) or the *Modulation Transfer Function* (MTF). The PSF represents the image of a single (infinitely small) point in an imaged object. The MTF is the normalized (maximum scaled to 1.0) Fourier spectrum of the PSF. The MTF describes the progressive loss of image contrast due to blurring as spatial frequency increases.
- Both PSF and MTF may be difficult to measure directly in an imaging system due to the difficulty of producing suitable test objects. A 1D approximation to the PSF can be derived from images of lines or edges which give, respectively, the *Line Spread Function* (LSF) and *Edge Spread Function* (ESF).