

Chapter 6

Contrast Adjustment

6.1 Introduction

The goal of all the medical imaging methods is to record a signal with sufficient resolution in intensity, space, and possibly time, to enable a medical interpretation of tissue structure and function. The information of interest, contrast of some sort, should be recorded in the intensity data – in other words, there should be sufficient intensity resolution in the raw data to discriminate the contrast details we are interested in. For a simple X-ray image on film this would mean that we have not under or overexposed the image. If a digital image constructed from the raw data is to be interpreted visually by a human, then we want the contrast information to be easily visible. To achieve this we will usually have to adjust the contrast of the image. This can be done in general terms to ensure that the full range of raw intensity information is distributed across the full range of display intensities, and in specific terms to exaggerate the displayed intensity differences for the part of the raw data intensities that represent the tissue properties of interest.

Contrast adjustment is a *point operation*, by which we mean adjustments are made to the intensity or color properties of pixels without regard to the properties of their neighbors. In a point operation there is no change in image size or geometry. A contrast adjustment *remaps* the original pixel intensities or colors into a new range of intensities or colors using a single *mapping function* that is applied to all pixels. The mapping function is a mathematical expression that describes how each pixel's intensity is changed according to its value in the original image. When used in this context 'mapping' does not imply any spatial rearrangement of intensity information.

6.2 Human Visual Perception

Most medical images are interpreted visually by humans so we have to produce and process these images in a form that is optimized for human vision. The human visual system can *adapt* to light intensities ranging over an amazing ten orders of

magnitude. However, at any instant the range of intensities (photon flux rates) that can be discriminated, that is, they are perceived somewhere on the scale between black and white, is only about four orders of magnitude (1–10,000). Within this range only about 24–30 distinct brightness *differences* can be perceived reliably. Small changes in intermediate light intensities are more easily perceived than the same size small changes in low and high intensities. This does not mean that we only ever need to display images with 24–30 distinct intensity levels. Visual perception adapts very quickly so that, as our eyes scan around an image, any local region of the image is perceived according to the local range of brightness. If we want to see fine details in an image we have to get very close to it – this improves our perception of both spatial *and* intensity details.

As well as intensity resolution we need to consider how the spatial resolution of human vision affects the images we create and view. In the region of highest visual acuity the eye has about 400 sensor cells per millimeter of retina surface. When an image is viewed with the eyes about 40 cm from the image surface each centimetre of image is focused onto about 0.4 mm of retina, or about $400 \times 0.4 = 160$ sensor cells. We would thus need an image resolution of at least 160 pixels/cm (or 400 DPI) to ensure that individual image pixels are not visually obvious. In reality the individual sensor cells in the retina are not fully independent (there are nerve cells interconnecting adjacent sensor cells) and the actual resolution limit of human vision is about 120 pixels/cm (300 DPI) at normal ‘close up’ picture viewing distances.

The spatial resolution of human vision can be expressed independently of viewing distance in terms of ‘cycles per degree’. In this case ‘cycles’ refers to spatial frequency and ‘degrees’ to an angle in the visual field. The maximum detectable spatial frequency is about 60 cycles/degree, while the maximum *sensitivity* occurs at about 3 cycles/degree (Fig. 5.3).

6.3 Histograms

A simple way to assess the overall contrast of an image, or more specifically, the range of intensities in the image matrix, is to plot a histogram of the intensities (Fig. 6.1). The histogram is a plot of the number of pixels (*y*-axis) of each intensity value (*x*-axis). The *y*-axis is often labeled ‘frequency’ in histogram plots. This use of the term frequency refers to the frequency of occurrence of an intensity value in the image. It is NOT in any way related to *spatial frequency* and the two usages should not be confused.

In Fig. 6.1, an MRI image of a lemon, the histogram shows a very large number of black (intensity = 0) and very dark pixels corresponding to the background, which is air surrounding the lemon, and the pith which contains very little water. The MRI signal comes mainly from water, so only system noise is present in the background of the image. There are much smaller numbers of gray and white pixels representing the juicy parts of the lemon in which the water gives a relatively intense MR signal.




Fig. 6.1 An MR image of an intact lemon and its corresponding histogram. The histogram is plotted on both linear (*black*) and logarithmic (*gray*) scales. The *y*-axis represents the frequency of occurrence of individual pixel intensities. Note the relatively small number of very bright pixels in the image and the corresponding low pixel count at the right end of the histogram

Because a large count of some intensities may make it difficult to assess the count of other less numerous intensity values in a histogram, it sometimes helps to plot the log of the pixel count. The log plot is shown superimposed in gray in Fig. 6.1.

In ImageJ use Menu: Analyze >Histogram to plot a histogram of the current image or current selection within an image. The ‘Log’ button adds in gray a supplementary plot of the log of the pixel count. Pressing ‘Alt-H’ on the keyboard enables plotting of a binned histogram as shown in Fig. 6.2.

A standard histogram plots one vertical bar for each possible intensity level in the image according to the bit depth of the pixel intensity data. To simplify the histogram and eliminate spikes and gaps it is sometimes useful to group intensity ranges into ‘bins’ and to plot one vertical bar for the total pixel count in each bin. This is termed a *binned histogram* (Fig. 6.2). Binning is also used when displaying histograms of images with a very high bit depth.

In a histogram *all* pixels with the same intensity are counted, irrespective of their position in the image, and the total count for each intensity level, or bin, is plotted as a single vertical bar. The histogram thus contains *no information* about the spatial distribution of pixels within the image. It is quite possible for two visually very different images to have identical histograms (Figs. 6.1 and 6.3, for example). Due to the lack of spatial information a histogram *cannot* be used to assess the relative contrast of objects or regions in an image. A histogram can, however, be used as part of a process that will result in the enhancement of the contrast of image objects.

A common use for histograms is to display the *dynamic range* of intensity data. In digital cameras a histogram can be used to check for *underexposure* and *overexposure* during both manual and automatic exposures. Simply reviewing a captured




Fig. 6.2 These are 'binned' versions of the histogram of the lemon MR image (Fig. 6.1). The user specifies the total number of bins. The 'Bin Width' refers to the number of adjacent intensity levels that are combined into each binned count (Bin width = total number of possible intensity levels divided by number of bins)

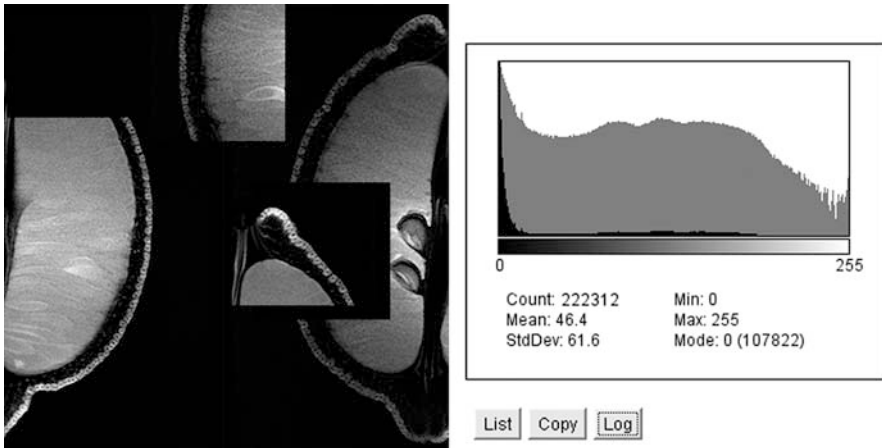


Fig. 6.3 The lemon MR image (Fig. 6.1) sliced up and rearranged. Note that the spatial rearrangement of the intensity data has no effect on the image histogram because the histogram simply displays the number of pixels of each intensity *irrespective* of pixel position. The histogram thus contains *no spatial information* and tells us nothing about the visibility of objects in the image

image on a camera's LCD display is a very unreliable way to check exposure due to the limited spatial and brightness resolution of this display. The light metering system of a digital camera measures the average light intensity incident on the sensor, or, in more sophisticated cameras, particular regions of the sensor. For automatic exposures the camera uses these light intensity measurements to set exposure parameters (aperture, shutter speed, and sensor signal amplification) to obtain image data that are calculated to utilize the full dynamic range of the sensor and the internal image processor. However, this approach will fail if there is a very bright region in the frame that was not measured, or if the majority of the measured region is dark and a small part is very bright. In this case the image data for the bright region

will be overexposed – all the pixels in the region will have the maximum possible intensity and any spatial variations in intensity within the bright region will be lost.

Just as failure to consider bright regions of the visual field can result in overexposure, the opposite problem, underexposure, can occur when exposure calculations include irrelevant bright regions. In this case subtle intensity variations in a predominantly dark subject will be compressed into a fraction of the available dynamic range and the signal-to-noise ratio will be reduced.

In manual exposure adjustment the photographer can use the histogram to check for appropriate utilization of the dynamic range *according to the particular subject and the information of interest within that subject*. This is an important point in imaging in general. We are nearly always making subjective judgements and predictions about the relevant information within the field we sample (the Field of View of a medical imaging system, or the frame of a camera exposure) and within the acquired image data. We have to make sure that the imaging device captures the information of actual or potential interest and, once captured or recorded, we have to extract that information from the image data.

6.4 Manual Contrast Adjustment

6.4.1 Contrast Stretching

Figure 6.4 shows a low contrast version of Fig. 6.1. This image is instantly recognizable as being of lower quality than Fig. 6.1. It is much more difficult to see the structural detail in the tissue of the fruit. The histogram, and our eyes, tell us that the darkest areas of the image are not very black, and the lightest areas not very white. The *dynamic range* of the data, 54–192 according to the Histogram tool, is much less than the dynamic range of the system which is 0–255. Notice that the overall profile of the histogram is the same as in Fig. 6.1, i.e. the relative numbers of dark and bright pixels is unchanged. We can predict that more information about the tissue structure is available than we can actually perceive in this image. This information may be perceptible if we can *expand* the current range of displayed intensities.

We can improve the contrast of Fig. 6.4 by expanding the dynamic range of its gray scale to the full dynamic range of the display system using the ImageJ Brightness/Contrast tool. This tool displays a miniature version of the histogram inside a box that contains a diagonal line (Fig. 6.5). Notice the changed position of this diagonal in Fig. 6.5b. The diagonal line represents a *brightness mapping function*.

A mapping function describes how the range of pixel intensities in the input image data is ‘remapped’ to a new range of intensities in the output image. The simplest mapping functions are *linear*. Figure 6.6 is an example of a linear mapping function that expands the central part of the range of possible input intensities to




Fig. 6.4 A low contrast image and corresponding histogram. The histogram indicates the absence of both high and low intensities in the image data – only about 60% of the full *dynamic range* of the display system is utilized

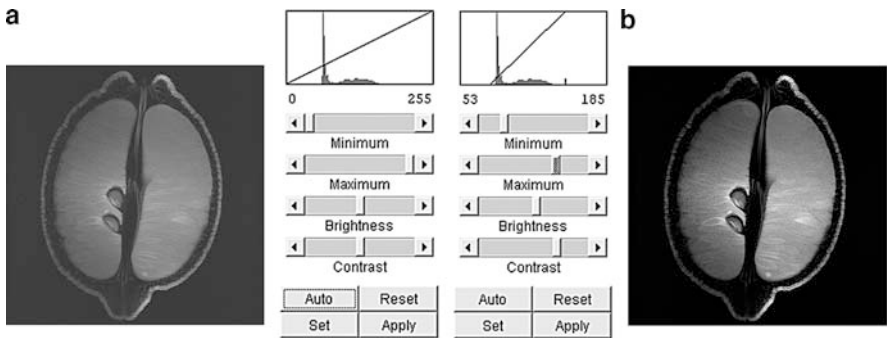


Fig. 6.5 Adjustment of image contrast using the ImageJ tool Menu: Image>Adjust>Brightness/Contrast. **(a)** Original low contrast image. **(b)** Improved contrast resulting from adjustment of the Minimum and Maximum sliders in the Brightness/Contrast tool. Note that the contrast adjustment tool window for image **b** shows the ‘input’ histogram *prior* to adjustment

fill all of the possible output intensity range. In this example only a fraction of the input dynamic range is being utilized and it is this range that we expand. It would be quite possible to adjust the mapping function so that a broader range of input intensities was used (by decreasing the slope of the mapping function) with the result, for this example, that the output dynamic range would not be fully utilized and the output image contrast would be sub-optimal. Alternatively, by making the slope of the mapping function steeper, only a fraction of the actual range of input intensities would be expanded. Input intensities below the minimum of the mapping function diagonal line would all be mapped to zero (= black) and any above the top of the diagonal would be mapped to 255 (= white). The mapping of the extremities of the input value range to the minimum and maximum possible output intensities is

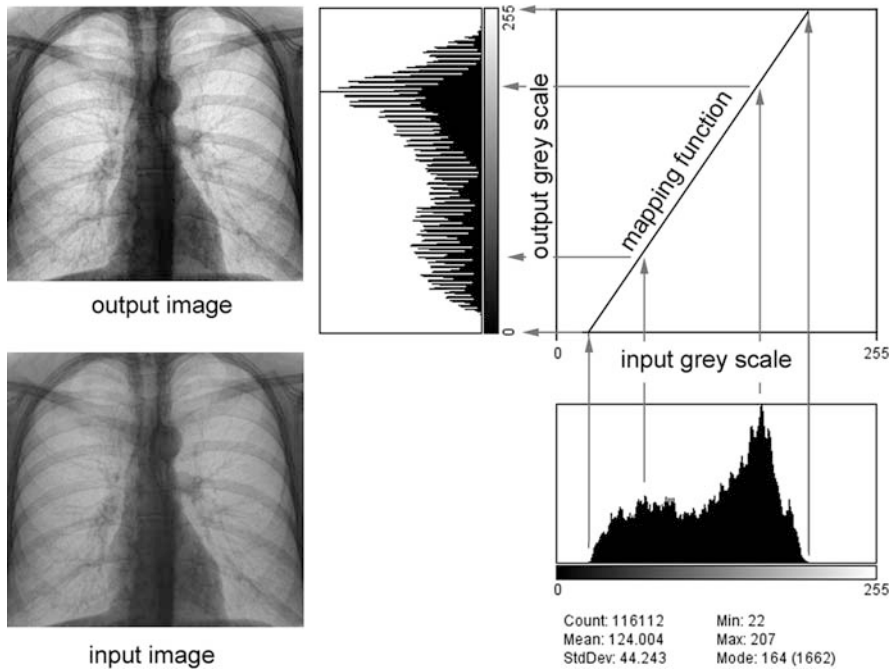


Fig. 6.6 Adjustment of image brightness and contrast by remapping of intensity levels. In this example a *linear mapping function* is used. A selected range of intensities (22–207) has been expanded to fill the available dynamic range (0–255). The histograms are shown as summaries of the range of input and output intensities – they are not essential to the process of contrast adjustment

termed *saturation*. It is often performed intentionally as such extreme values usually represent noise rather than information.

Expanding the range of image intensities is referred to as *histogram stretching* but this term is misleading. The mapping function is NOT applied to the input histogram. It is applied to the input image data of which the input histogram is just a summary. Histogram stretching is thus a *result* of the remapping of the intensity range. It is not an essential part of the process. We can perform contrast expansions using mapping functions without ever calculating or viewing input or output histograms. We might perhaps use the input image histogram as an aid to choose the minimum and maximum values for the mapping function and this is the reason a miniature input image histogram is displayed within the ImageJ 'Brightness/Contrast' tool.

A linear contrast expansion will typically give rise to a histogram that contains regularly spaced gaps – intensities for which the pixel count is zero. Conversely, a linear contrast decrease causes regularly spaced spikes to appear – intensities with unusually high pixel counts. Both these phenomena are the result of *quantization* of the intensity values. Because adjusted intensity values must be rounded before they are stored as integers, specific integer intensity values may end up with extra or zero




Fig. 6.7 The origin of spikes and gaps in histograms following contrast reduction and expansion. In this simplified example the original intensities (0–31, bit depth = 5) are either reduced or increased by 10% – a linear contrast adjustment. The adjusted intensity values must be rounded to the nearest integer as only the integers 0–31 can be stored in this example. After a 10% contrast decrease rounding results in the doubling of the number of pixels with intensities 5, 14, and 23. These intensities appear as spikes in the new histogram. Conversely, after a 10% contrast increase rounding results in the absence of pixels with intensities 5, 16, and 27. These intensities appear as gaps in the new histogram. Non-linear mapping functions may lead to spikes and gaps in different regions of the histogram of the adjusted image (e.g. Fig. 6.10)

pixel counts (Fig. 6.7). The presence of spikes or gaps, or both, in a histogram is an indication that the image has previously undergone some post-acquisition contrast adjustment.

What is the difference between a mapping function and a lookup table? Both describe the range of output intensities that should be substituted for input image data. In general a mapping function is a mathematical expression used to *permanently*




Fig. 6.8 The ImageJ ‘Show LUT’ (Show Lookup Table) command displays a color plot of the RGB values in the lookup table data. This is NOT a conventional mapping function display. This example illustrates the data from the ‘Union Jack’ lookup table that was applied to produce Fig. 2.9c. Note that both *white* (255) and *mid-gray* (127) inputs produce a white output from this lookup table – thus some of the original intensity information is lost in a display that uses this lookup table

change the intensity data in an image matrix or image file. In contrast, a lookup table is a table that describes how to *display* image data. Application of a lookup table, or changing the lookup table of an image file, does not change the stored intensity data.

In ImageJ the Menu: Image > Color > Show LUT command (Show Lookup Table) displays a color plot of the RGB values in the lookup table data (Fig. 6.8). The display is very similar in appearance to a mapping function because the lookup table is ‘mapping’ each original intensity to a specific display color. The ‘List’ button displays the actual lookup table.

As mentioned above, a simple expansion of the histogram to match the dynamic range of the display system might improve overall contrast but may not provide optimal enhancement of the information of interest. Remember that while the histogram gives a useful summary of overall image contrast it does not tell us much about the ease of interpretation of the image. To optimize the contrast of some specific image feature (e.g. anatomy) it is generally necessary to adjust the contrast while looking at the image rather than the histogram.

In ImageJ use Menu: Image > Adjust > Brightness/Contrast to perform a linear adjustment of image contrast. This tool provides four controls to adjust only two parameters. It is easiest to make adjustments using either just the top two sliders, or just the bottom two.

6.4.2 Window and Level

In medical imaging, especially CT, it is common to refer to linear contrast adjustments with the terms *Window* and *Level* (Fig. 2.4). In this context the *Window* (or Window Width) describes the width of the range of image data intensities that will be expanded to fill the dynamic range of the display. This corresponds to the horizontal ‘span’ of the diagonal part of the linear mapping function. The *Level* (Window Level) specifies the center of the Window range. All pixels outside the window are saturated – changed to either black or white. Figure 6.9 illustrates the use of the ImageJ ‘Window&Level’ tool to selectively enhance the contrast of either bone or metal implants in a projection X-ray image.

In CT the image data represents calculated CT numbers rather than arbitrary intensity units. Since the CT number is directly related to the linear attenuation coefficient of the imaged tissue, defined windows and levels are used to optimize the display of contrast for particular tissue types. Examples of bone and soft tissue windows were illustrated in Fig. 2.5.

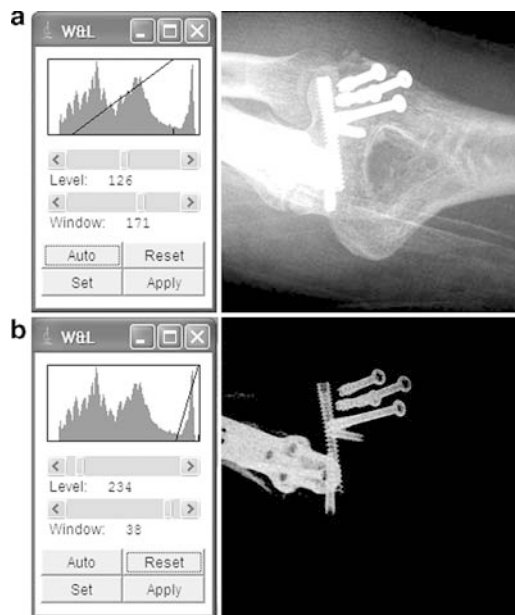


Fig. 6.9 Selective linear contrast adjustment using Window and Level. The *Window* describes the width of the range of image data intensities that will be expanded to fill the dynamic range of the display. The *Level* specifies the center of the Window range (see also Fig. 2.5). In this example window **a** is selected to optimize the contrast for bone and soft tissue. Window **b** eliminates all tissue and highlights the detail of the surgical metal implants

In ImageJ the Menu: Image > Adjust > Window/Level tool behaves identically to the Contrast and Brightness slider controls in the 'Brightness/Contrast' tool.

6.4.3 Nonlinear Mapping Functions

A linear contrast expansion is ideal for distributing the intensity data of interest across the full dynamic range of a display system without changing relative intensities within that range. However, we may also want to enhance the contrast of images (or regions of images) which already utilize the full dynamic range available. We want to enhance differences in intensity without saturation of the brightest or darkest intensities. Put another way, we want to make the dark regions darker, and the light regions lighter without losing intensity information due to saturation. This contrast adjustment can be performed with a non-linear mapping function (Fig. 6.10).

6.5 Automatic Contrast Adjustment

Why bother with a histogram at all? So far we have used histograms for nothing more sophisticated than finding the minimum and maximum intensities in an image and checking that the acquired data lies appropriately within the dynamic range of the system. The histogram is much more useful in automated contrast adjustment.

6.5.1 Normalization

The contrast stretching operation just described using the Brightness/Contrast or Window/Level tools can also be performed automatically if we specify the type of adjustment we want to make. When a linear mapping function is used to stretch the dynamic range of the input image to fill the available dynamic range of the storage or display system the operation is called *Normalization* (or, somewhat inappropriately, 'Histogram Normalization').

The simplest form of normalization stretches the histogram as we did above to set the input minimum intensity to zero in the output and the input maximum to the available maximum (255 if bit depth is 8). The problem with this method is that in a predominantly low contrast gray image a single pixel with a very high or low value will prevent the full enhancement of the image. We can get around this problem by assuming that the information of interest does not include extreme intensities and




Fig. 6.10 A non-linear mapping function can be used to enhance contrast without saturation of the highest or lowest intensities. In this example a sigmoid shaped mapping function is used to darken all low intensities and brighten all high intensities. The effect is to enhance the visibility of bone details without losing all soft tissue detail. Note the presence of gaps in the center of the output image histogram, due to contrast expansion of the middle range of intensities, and spikes in the histogram at low and high intensities, due to contrast compression of these intensities (described in Fig. 6.7). The illustrated mapping function is from Adobe Photoshop. Non-linear manual contrast adjustment is not currently available in ImageJ

permit a small percentage of pixel values (the extreme ones) to saturate. The linear contrast stretch will then be performed over a range slightly less than the full input range.

In ImageJ use Menu: Process > Enhance Contrast to perform automatic normalization with adjustable saturation.

Figures 6.11 and 6.12 illustrate the use of normalization to improve contrast in an MR image. The percentage of total pixels (dark and bright) deemed to be extreme has been set in the ‘Saturated Pixels’ box of the ImageJ ‘Enhance Contrast’ tool. As the Saturated Pixels setting is increased the peak of the original histogram




Fig. 6.11 Image normalization with variable saturation. (a) Original low contrast MR image of an orange. (b) Normalized image with Saturated Pixels = 1%. (c) 2%. (d) 8%. The percentage of total pixels, dark and bright together, was set in the ‘Saturated Pixels’ box of the ImageJ ‘Enhance Contrast’ tool. Note that in the extreme case, when the percentage of pixels excluded from the normalization is too large, image detail is lost in the brightest regions. The histograms of these four images are shown in Fig. 6.12

broadens and shifts towards the high intensity end of the scale, consistent with the overall increase in brightness of the image. Note that in the extreme case, when the percentage of pixels excluded from the normalization is too large, image detail is lost in the brightest regions.

6.5.2 Histogram Equalization

If we temporarily ignore the fact that histograms tell us nothing about the spatial distribution of intensity information in an image and think about images *only* in terms




Fig. 6.12 Histogram normalization using increasing ‘Saturated Pixels’ settings (0, 1, 2, and 8%). The histograms are those of the orange MR images shown in Fig. 6.11. As the ‘Saturated Pixels’ setting is increased the central peak of the original histogram broadens and shifts towards the high intensity end of the scale, consistent with the overall increase in brightness of the image

of their histograms, then we might reasonably assume that an image with ‘ideal’ contrast properties would have its intensity information evenly distributed across the available range. Such an image would have a histogram with a flat profile – there would be roughly the same count of pixels for each possible intensity. *Histogram Equalization* is an automated process that does just this – it attempts to redistribute an image’s pixel intensities to produce a histogram with equal counts for all possible intensities.

From the regular histogram we can plot a *cumulative histogram* which is a running sum of the total pixel count starting from the lowest possible intensities and progressing to the highest (Fig. 6.13). The cumulative histogram thus *never* decreases as we scan from left to right – its slope is always positive or zero. When an image contains no pixels in a particular intensity range – a gap in the regular histogram – the cumulative histogram will be flat in that range. When there are a lot of pixels in a particular intensity range – a peak in the regular histogram – the cumulative histogram will increase rapidly.

You may notice from Fig. 6.13 that the plot of the cumulative histogram looks similar to a mapping function. What would happen if we used the shape of the cumulative histogram as a mapping function to adjust the contrast of an image? This is exactly how equalization is implemented! The process is illustrated in Fig. 6.14. Notice, however, the raw histogram of the equalized image is not flat, as we should expect from equalization, but contains many spikes and gaps. This is because we are dealing with quantized intensity data. The spikes and gaps occur for the same reason we saw them after linear contrast adjustments (Fig. 6.7). To see the expected flat histogram we need to bin the intensities into ranges. The use of the cumulative histogram to describe a mapping function that will flatten an image’s histogram is more obviously demonstrated with a simple image such as that shown in Fig. 6.15.

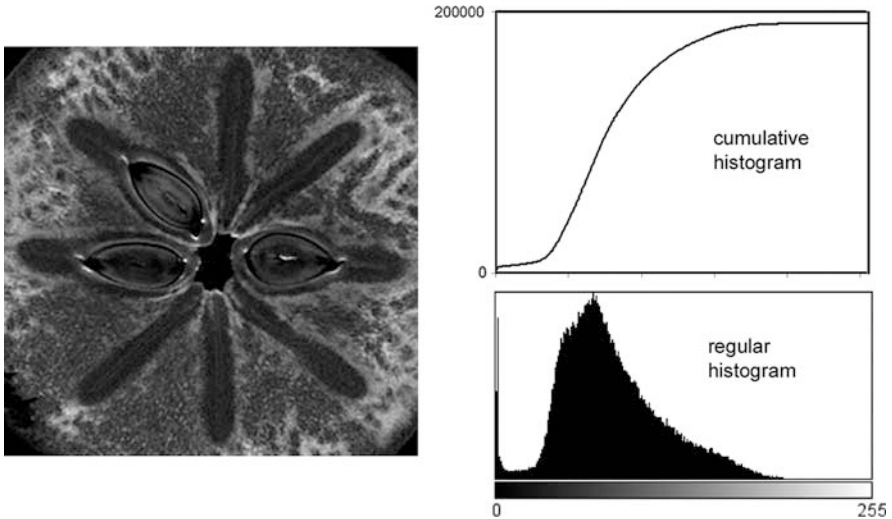


Fig. 6.13 The cumulative histogram represents a ‘running sum’ of the pixel counts in the regular histogram or, equivalently, the total number of pixels with an intensity less than or equal to a specific intensity. Notice that the cumulative histogram is flat in intensity ranges where there are low or zero pixel counts and increases rapidly in ranges that have high pixel counts

Because equalization is directly based on an image’s histogram, and the spatial information in the image is absent from the histogram, the visual effects of equalization depend very strongly on the individual image. For most images it turns out that equalization using a mapping function based on the simple cumulative histogram produces an image that appears, to a human, to be too bright and lacking in overall contrast detail. To reduce this common adverse result a mapping function based on the *square root* of the pixel counts can be used (Fig. 6.16). This is the default method in ImageJ. Figure 6.17 compares the effects of normalization and equalization applied to a low contrast X-ray image.

Equalization is an option in the ImageJ ‘Enhance Contrast’ tool. The default algorithm uses the square root of the pixel count from the original histogram to define the cumulative histogram and the mapping function. Hold down the ‘Alt’ key when pressing the ‘OK’ button to perform equalization using the raw pixel counts. If a ROI is defined within the image when the Enhance Contrast tool is opened then the option ‘Update All When Normalizing’ will be available. This option uses only the data inside the ROI for calculation of the mapping function, but applies the mapping function to *all* pixels in the image. In the case of equalization the ROI-based mapping function is always applied to the whole image.

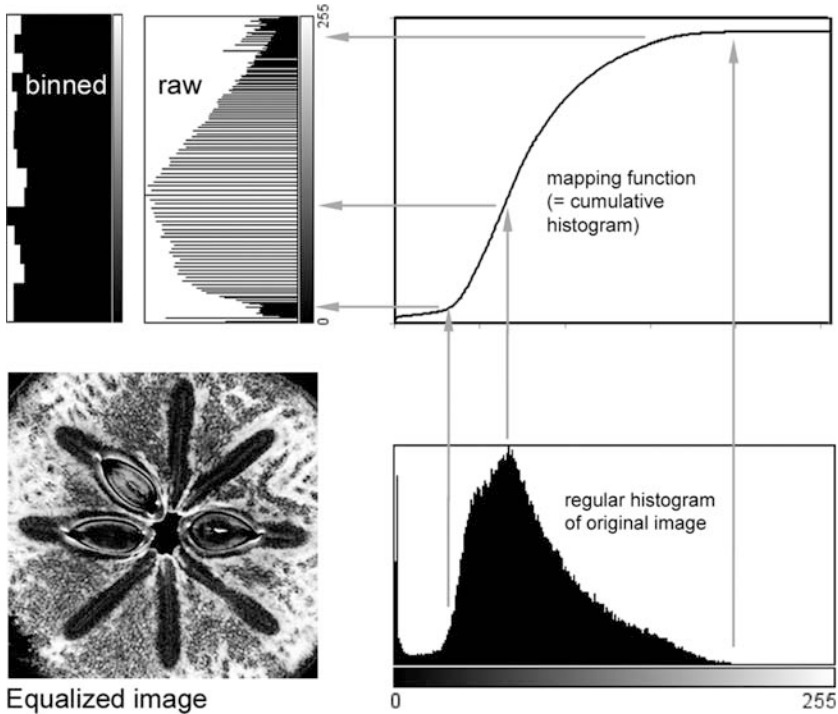


Fig. 6.14 Equalization of the low contrast MR image shown in Fig. 6.13. The equalization process uses the cumulative histogram to define the shape of the mapping function. The histogram of the equalized image has roughly the same number of pixels in all intensity ranges, however, due to quantization of the intensity data, the expected flat histogram is usually only seen after binning of intensities

6.5.3 Histogram Specification

We saw above that for many images equalization, or flattening, of the regular histogram does not produce visually ideal contrast properties. Most ‘good looking’ images actually have a roughly bell-shaped regular histogram. This is also the case for medical images. If we assume that some ‘ideal’ histogram shape exists, at least for a particular type of image, then perhaps we could adjust suboptimal images by making their histograms conform to this specific ideal. This can be done with a modification of the equalization process called *Histogram Specification*.

The equalization method uses the cumulative histogram of an image as the form of a mapping function that, in principle, flattens the regular histogram. We can, at least in theory, reverse this process – that is, take an equalized image with a flat histogram and use the mapping function backwards to get back to an image with a histogram the same as the original image before equalization. The ‘reverse’ mapping function is the inverse of the equalization mapping function. Now consider our image that has an ‘ideal’ histogram. If we were to equalize this image we would use

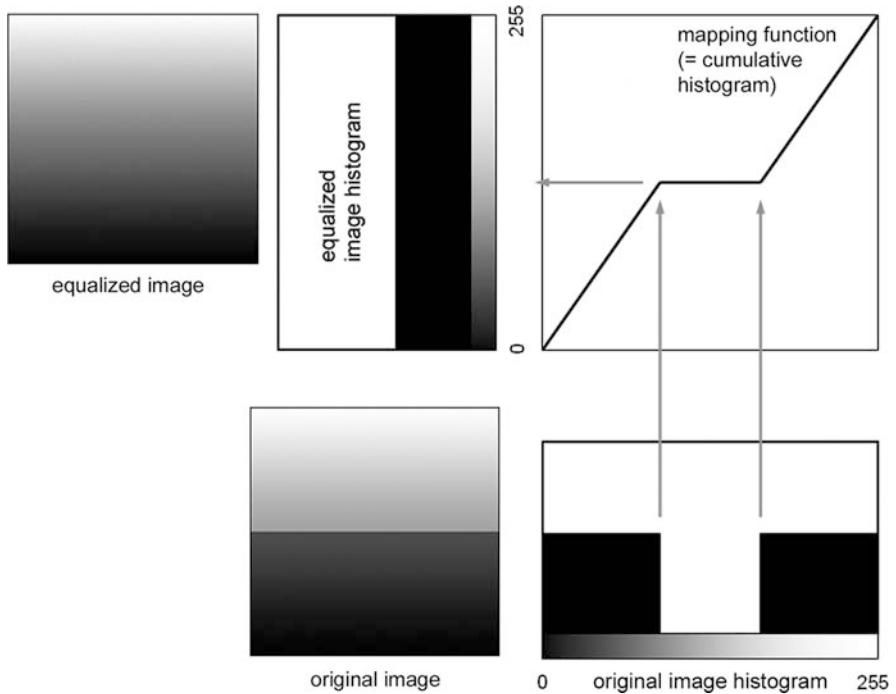


Fig. 6.15 A simplified demonstration of the equalization process. The cumulative histogram of the original image describes the form of a mapping function which will flatten the original histogram. The original image here has dark and light intensities but no pixels in the middle intensity range. Its regular histogram thus has a large central gap and its cumulative histogram a corresponding central flat region. The effect of using the cumulative histogram as a mapping function is to stretch the low and high intensity ranges such that they now meet in the middle and the gap in the original histogram is eliminated. The equalized image has an equal number of pixels at each possible intensity level

a mapping function that flattens its histogram. The inverse of this mapping function, let's call it the 'specification function', will change *any* image that has a flat histogram into one that has the 'ideal' histogram. Histogram specification is simply the two step process of first equalizing an image and then applying a specification function to 'unflatten' the histogram and make it the desired shape.

Histogram specification is a good idea in principle but suffers from the intensity quantization problem and only really works for images of quite similar spatial composition.

6.5.4 Region-Specific Contrast Adjustments

Automated normalization and equalization often fail to produce improvements in image contrast because irrelevant image information, for example background noise




Fig. 6.16 Comparison of equalization methods. (a) original low contrast image. (b) Equalization based on the raw pixel count. (c) Equalization based on the square root of the pixel count. For most images the square root method (c) produces a more visually attractive image, however, notice that the details of the persimmon seeds are easier to see in image **b**

or overlaid text, is included in the contrast adjustment calculations. This problem can often be avoided by defining within the image a Region Of Interest (ROI) on which to base the calculations. The mapping functions thus generated can be applied just to the ROI or to the whole image (see Section 6.6 for an example).

6.5.5 Binary Contrast Enhancement – Thresholding

Sometimes the only information we are interested in is the boundaries between regions in an image. We might for example want to select all the regions in an X-ray image that represent bone, or all the white matter from an MR image of the brain. The general name for this process – dividing an image into distinct regions – is *Segmentation*. Under ideal conditions distinct anatomical regions will be associated with unique ranges of intensities. Better still, we might be able to say that there is a single cutoff intensity above or below which all pixels belong to a particular category, bone or not-bone, say in our X-ray image. In this ideal case we could define two distinct image regions simply by changing the intensity of any pixel that exceeds the threshold value to 255, or white, and changing any pixel less than or equal to the threshold to zero, or black. We would thus create a *binary image*, or *mask*, that contains only two possible pixel intensities. Binary images are typically shown in pure black and white, with the choice of which region is shown black and which white being somewhat arbitrary.



Fig. 6.17 Comparison of effects of normalization and equalization on contrast of an X-ray image. (a) Original image. (b) Normalized image (0% saturated pixels). (c) Equalization based on raw pixel counts overbrightens the image and exaggerates background noise. For this particular image equalization based on the square root of pixel counts enhances contrast detail in most of the anatomy, without too much noise exaggeration, however, some detail in the largest bone is lost

Imagine we are only interested in the bone details of an X-ray image such as Fig. 6.18. We can see that the bones are generally bright and select an appropriate threshold intensity (135 in this case) to create a binary ‘bone’ mask. A possible problem with this binary image is that it includes implant metal in the bone mask. Since most of the metal appears brighter than bone in the original image use of a *second* upper threshold enables the exclusion of most of the metal while retaining most of the bone.

When storing binary images we can save a lot of computer space by using the values zero and one, instead of zero and 255, to store the intensity information. For an eight bit image this represents an 88% saving since the required bit depth

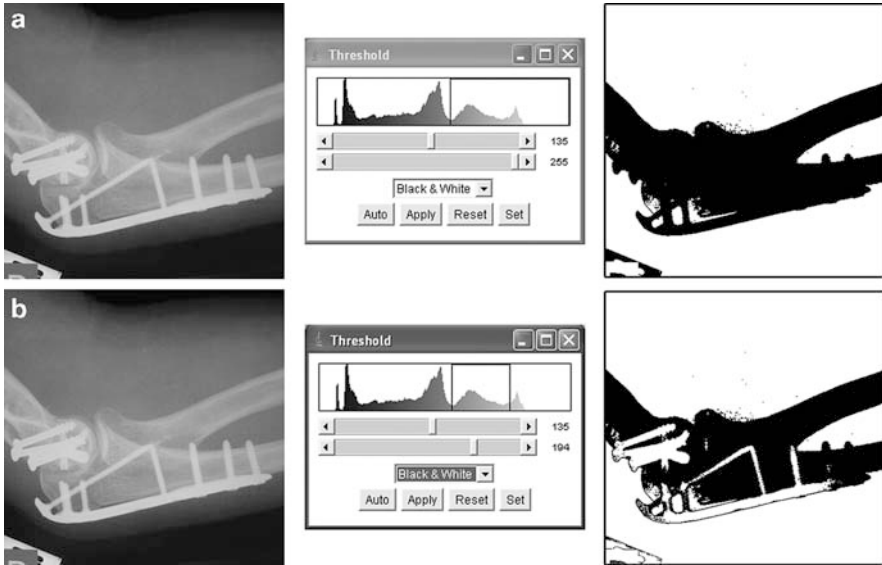


Fig. 6.18 Use of thresholding as a crude selection of the bone regions of an X-ray image. A single threshold intensity (a) results in inclusion of implant metal with bone. Using upper and lower threshold intensities (b) enables selection of bone without inclusion of metal. The method is imperfect – some bone is excluded and some metal is included. Thresholding is a primitive method of image segmentation




Fig. 6.19 A binary mask created by thresholding can be used to exclude irrelevant information from an image. Here the ‘bone’ mask from Fig. 6.18b is inverted and used to remove soft tissue and metal from the original image. Each original pixel intensity is multiplied by the corresponding mask pixel value (0 or 1). Image $c = a \times b$

is only one. Obviously a great deal of image information is lost in the thresholding process but this is quite intentional. In thresholding an image we are only interested in the question: ‘Which of two categories does a pixel belong to?’. Used as a mask a binary image can be multiplied pixel-wise by the original image to remove all the image information that is of no interest. For example, Fig. 6.19 illustrates the use of our ‘bone only’ mask (after inversion) to remove the soft tissue and metal implant details.

The ImageJ Menu: Image > Adjust > Threshold tool creates a binary image based on the selection of upper and lower thresholds. All pixel intensities above the upper threshold, or below the lower, are changed to the maximum possible (white). Pixels between the thresholds are changed to zero (black). This tool also enables a color display of the binary image. The tool Menu: Process > Image Calculator can be used to apply a mask to another image.

6.5.6 Hardware Contrast

Our discussions about image contrast adjustment made the implicit assumption that all image viewing devices will produce identical displays from identical image data. This assumption is rarely justified for gray scale images and is generally not true at all for color images. In medical imaging it is essential to check and calibrate image display devices to achieve consistent contrast visibility.

6.5.6.1 Gamma Adjustment

For a display monitor the output light intensity (I) is related to the source voltage (V) by a power law:

$$I \propto V^\gamma \quad (6.1)$$

Monitor gamma (γ) describes the adjustment that is made to image intensity data to achieve a *perceived* display intensity that is linearly related to the image data intensity. Figure 6.20 illustrates the effect of gamma adjustment. When $\gamma = 1$ no change is made to the raw intensity data with the result that the gray values appear too light. In Fig. 6.20 the gamma values of 1 or 1.2 produce a roughly linear gray scale gradient *in the printed image*. On a CRT monitor display, which is based on light emission rather than light reflection as is the case for this printed image, a gamma in the range 2.2–2.5 is usually appropriate. The light output of LCD monitors cannot be described by a simple power law as for CRT monitors. Proper adjustment of the output of LCDs requires calibration with a photometer and software that can modify the graphics card lookup table.

Some computer operating systems include a gamma adjustment tool. This tool is used to make small adjustments to gamma in order to make a patch of mid gray appear to have, when viewed at a distance, the same brightness as a grid of closely spaced black and white lines (Fig. 6.21). This adjustment sets the midpoint on the gamma curve to enable a roughly linear gray scale display but does not account for contrast differences due to luminance or overall brightness.

Image *acquisition* devices also have non-linear responses to input energy intensity that can be approximated by a power law. The important message here is that the

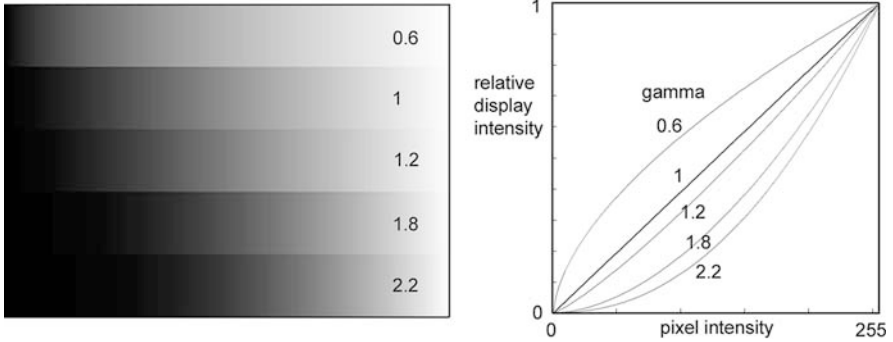



Fig. 6.20 Illustration of the effect of gamma adjustment where display intensity depends on input signal according to a power law: $I \propto V^\gamma$. When $\gamma = 1$ no change is made to the raw intensity data with the result that the gray values appear too light. Here the gamma values of 1 or 1.2 produce a roughly linear gray scale gradient in the printed image

Fig. 6.21 Part of the *Adobe Gamma* CRT monitor adjustment tool. The aim is to adjust gamma until the black/white bar pattern appears to have the same brightness as the central mid-gray region. This adjustment sets the midpoint on the gamma curve to enable a roughly linear gray scale display



signal intensities that are recorded as pixel values in an image file will, in general, NOT be linearly related to the energy emitted by the imaged objects, nor to the luminance of a display (or density of a print) created from the image file.

6.5.6.2 The DICOM Gray Scale Display Function

Gamma adjustment should not be confused with *monitor calibration* which is necessary to achieve consistent and predictable *perception* of image contrast differences. Monitor calibration requires a colorimeter device which is periodically placed on the face of the monitor to check and adjust calibration. Some of these devices can also automatically adjust the display intensity to compensate for changing room lighting conditions.

Medical image displays should be calibrated according to Part 14 of the DICOM Standard (www.nema.org) which defines a *Grayscale Standard Display Function* (GSDF). The purpose of the GSDF is to ensure that diverse medical image display

devices produce images of comparable quality in terms of speed and accuracy of diagnostic interpretation. Some newer medical monitors have built-in automatic GSDF calibration using an internal photometer.

6.6 Practical Example. Adjusting the Contrast of a Magnetic Resonance Microimage

Finally, let's take a look at a practical example of a contrast adjustment problem. In Fig. 6.22 we see a newly acquired MR image of an excised mouse brain immersed in formaldehyde solution in a cylindrical glass bottle. The image was acquired in an 11 Tesla microimaging system so the spatial resolution is much higher than could be achieved in a typical 1.5T or 3T human MR scanner. Nevertheless, this image illustrates a number of the problems commonly seen in clinical medical images. Most of the image background, aside from the text, is black – there is no MR signal detected outside the glass bottle because there is no water (more specifically, no hydrogen nuclei) outside the bottle. For the same reason we see no signal from the glass walls of the bottle. The detected signal arises from water protons in the brain tissue and in the formaldehyde solution outside the brain. The information of interest – the structure of the brain – is revealed by subtle contrast differences within the brain tissue, and the overall shape of the brain by the contrast difference between the brain tissue and the formaldehyde solution. The tissue structure information is squeezed into a range of gray scale intensities running from mid-gray to near white.




Fig. 6.22 High-resolution MR image of a mouse brain as typically displayed on the console of the imaging system, in this case a Bruker 11 Tesla system. The image uses the full dynamic range of the display system but within the brain tissue the contrast is relatively poor (image courtesy of Professor Bill Price, University of Western Sydney)

The scanner software overlays the image with text describing significant acquisition parameters. This text contains the only truly white pixels in the image.

Looking at the histogram we see a large Uluru-like hump in the middle range of intensities. It is tempting to assume that this hump represents the brain tissue, but such an assumption might be a mistake. Remember that the histogram is simply a display of the relative number of pixels of each intensity – it tells us nothing about the spatial distribution of those pixels.

In order to improve our ability to see the structure of the brain we need to adjust the contrast of the image. In particular, we want to *expand* the contrast of the region of interest, the brain tissue, and we don't particularly care about what happens to the contrast outside the brain tissue, except perhaps we don't want to lose the readability of the text. We can make a manual contrast adjustment with the ImageJ Brightness/Contrast tool, as shown in Fig. 6.23. We increase the Minimum slider to make the darkest parts of the brain almost black, and decrease the maximum slider to make the brightest parts of the brain almost white. It is important that we make these adjustments while looking at the image and not just the histogram. We want to enhance the contrast without losing structural information. If, for example, we set the Maximum slider a little too low, some of the anatomical detail in the brightest regions of the brain is lost.

We have just discussed a method of making a *linear* adjustment to the image contrast based on our visual assessment of changes in the region of interest as we manually adjust the contrast – a linear remapping of the grayscale. In some images we could achieve the same improvement in contrast, or possibly improve on it, if we used an automated contrast adjustment method such as histogram equalization. Figure 6.24 illustrates the output image after we equalize *the whole image*. The result is unsatisfactory in terms of visualization of the brain structure because the equalization algorithm was based on the cumulative histogram of the whole image. The pixels forming the white overlaid text and the dark background have been included in the calculation. In fact the noise in the background has been exaggerated. It would be preferable to perform the equalization based on a histogram that represents only the pixels in the region of interest. We can do this by manually defining the ROI.

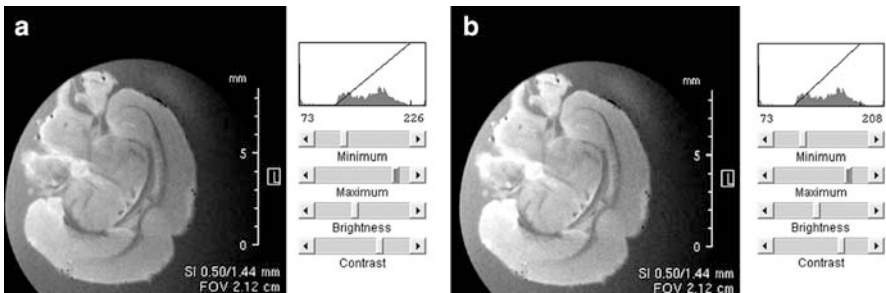


Fig. 6.23 Linear contrast enhancement of the brain tissue in Fig. 6.22 using the ImageJ ‘Brightness/Contrast’ tool. In image **b** the ‘Maximum’ slider has been set too low. Note the loss of tissue structure details in the bright regions. In image **a** the enhancement is close to ideal




Fig. 6.24 Contrast adjustment based on equalization of the histogram of the whole image, including text and background. Note the noise now visible in the background and the lack of contrast enhancement in the brain tissue




Fig. 6.25 Manual selection of a Region Of Interest (ROI), and display of the histogram of the ROI, permits exclusion of irrelevant pixel intensities from the contrast adjustment process

ImageJ provides several tools for selection of specific regions of an image on which subsequent operations or calculations will be performed. In Fig. 6.25 we have used the freehand selection tool to define a ROI that just includes the mouse brain. When a ROI is currently defined within an image the Histogram tool displays only the data from within the ROI. Note the similarity between the shape of this histogram and the central hump of the histogram of the whole image (Fig. 6.22). In this particular image our earlier warning about making assumptions about the spatial origin of a histogram feature turned out to be over-cautious. Now the mapping function for the equalization will be based on just the pixel data within the ROI. The equalization can then be applied to all the pixels in the image (Fig. 6.26).




Fig. 6.26 Contrast adjustment of the whole image based on equalization of the histogram of the ROI defined in Fig. 6.25. Now there is good tissue contrast. The loss of contrast details outside the ROI is of no consequence

In this example we manually defined the ROI. We made a subjective expert judgment about which part of the image contained the important information and we applied user input to constrain the automated calculations and adjustments. It seems trivial for a human inspector of the mouse brain MR image to point to the parts of the image that represent the brain tissue, but this task is far from trivial for a machine which has to somehow know not just what a brain looks like but also what is noise or text added after acquisition of the image.

Subjective expert input is the norm in medical imaging because human bodies, and biological structures in general, are extremely diverse and difficult to characterize in simple geometric ways that are amenable to automated machine analysis. Only in a very few specific contexts can a machine more reliably recognize pathology than an expert human.

6.7 Summary

- Image data is usually measured and stored with greater precision than the intensity resolution of human perception. Contrast adjustment is the modification (remapping) of stored pixel intensities in order to make specific intensity information more visually obvious in the image display.
- A histogram provides a summary of the range and frequency of occurrence of pixel intensities in a digital image. However, the histogram provides *no* information about the spatial distribution of pixel intensities within the image. Two completely different images may have identical histograms.

- *Linear* contrast adjustments use a linear mapping function to expand a range of input pixel intensities. *Normalization* is the process of expanding the *whole* input intensity range to the full available dynamic range.
- A *fraction* of the input range may be expanded with a linear mapping function to optimize the display of particular image features. In CT, for example, defined *Windows* are used to display specific ranges of the stored CT number data using the full dynamic range of the display device.
- *Equalization* uses an image's *Cumulative Histogram* to describe the form of a non-linear intensity mapping function. The effect is, theoretically, to flatten the regular histogram and thus produce an image with the same number of pixels of each possible intensity. Because intensity data is quantized the equalized histogram usually has many spikes and gaps, however, a *binned histogram* of the image will be approximately flat.
- *Thresholding* is an extreme form of contrast adjustment. All pixels intensities within a specified range are remapped to the available maximum intensity, and all other pixels are remapped to the available minimum intensity. Thresholding converts a gray scale image to a pure black and white (binary) image.
- Imaging detectors typically have a non-linear response to input energy. Similarly, image display devices typically have a non-linear response to input data. The response of many devices can be described by a power law and can be adjusted by a single parameter – *gamma*.