

Chapter 6

Digital Imaging



W. Gray (Jay) Jerome

6.1 Introduction

Image capture and rendering in confocal microscopy is a digital technique. To optimize confocal image acquisition, the operator needs to have a basic understanding of what constitutes a digital image and how digital imaging can be employed effectively to help visualize specific details. In this chapter, we will introduce the topic of digital images, discuss the basic components of the digital image, and provide some basics on how to optimize collection of these components with regard to confocal microscopy.

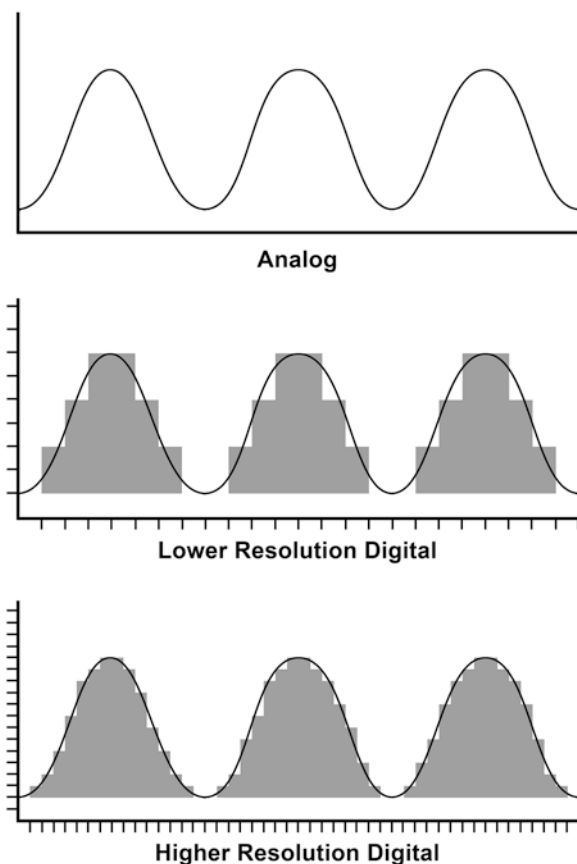
6.2 Analog Versus Digital Information

An analog signal is a continuous one. Its possible values can be $-\infty$ to $+\infty$ in infinitely small increments. This means that values such as 1, 2, 2.02, and 3.0000023 are all possible. In fact, the set of possible values includes all real numbers. In comparison, a digital signal is composed of discrete elements that have values of one unit or some whole multiple of one unit. Thus, you cannot have a value that is a fraction of one unit, such as 0.33, or a fraction of a whole multiple, such as 2.57. Of course, at this point we have not defined the basis of our unit. In fact it could be a micrometer, a yard, or 2.65 feet. However, once the unit is defined, then all elements of the set must be in whole multiples of that unit, and no element can be less than one unit in size. Thus, the value of 1 unit (1 quantum of information) defines the lower limit of resolution.

W. G. Jerome (✉)
Department of Pathology, Microbiology and Immunology,
Vanderbilt University School of Medicine, Nashville, TN, USA
e-mail: Jay.Jerome@Vanderbilt.edu

Analog to digital conversion consists of converting the set of analog values to a set of digital values. Two steps are involved in this conversion. The first is determining the value of one unit. The second step involves deciding the means by which the analog values are assigned to one of the discrete digital values. This process is depicted in Fig. 6.1 for analog values varying with respect to an X and a Y parameter. The value of 1 quantum in the X and Y directions is indicated by the width of the tic marks. The Y value for a specific quantum on the X-axis is determined by taking all the analog values occurring within the range of a single quantum on the X-axis and converting them to the nearest quantum along the Y-axis. Often this is done by taking the average of all the values within one X quantum and then rounding to the nearest Y quantum value. In the low-resolution conversion of Fig. 6.1, a relatively high value has been chosen for 1 quantum of X and Y. In the high-resolution conversion, a smaller value has been chosen. Of course, this smaller value results in a closer approximation of the digital graph to the original analog information. Similarly, selecting a high-resolution value in the confocal microscope software will result in a better approximation of the analog information viewed

Fig. 6.1 Analog to digital conversion of a linear signal





through the microscope with our eyes. However, as discussed below and in other chapters of this book, there are guidelines in selecting the appropriate resolution values available in the software. It is not always the best practice to select the highest digital resolution available.

6.3 Pixels and Pixel Size (Spatial Resolution)

In the case of a two-dimensional micrograph, the spatial information in the image is divided up into a series of discrete two-dimensional picture elements (pixels). A pixel is one unit wide by one unit high. A pixel does not have to be square; in other words one quantum in width does not necessarily have to be the same as 1 quantum in height. However, in microscopy, square pixels are usually employed, and we will confine our discussion for the rest of this chapter to square pixels.

A two-dimensional digital image is made up of a mosaic of pixels. Besides the determination of the X and Y quanta, a third value must be determined, that of what value to paint in the pixel. Whatever the value, it will be uniform within the pixel. For a fluorescence image, we must convert the analog brightness information in the original sample to a digital representation. To do this, all of the values within a pixel area are averaged to a single value as depicted in Fig. 6.2. For a gray-scale image,

Fig. 6.2 Analog to digital pixel conversion



Averaging Gray
Levels to
Determine Gray
Value for Pixel




Fig. 6.3 Effects of different pixel sizes on image resolution

the value would be some gray value between black and white. In the simplest conversion, this would involve using the average of all the gray values in the pixel. For example, a pixel that contained 50% black and 50% white would be averaged to a gray value at the middle of the scale. Because all of the details within a pixel have been converted to a single value, the smallest detail that can be represented in the digital image is no smaller than one pixel.¹

It follows, then, that the smaller the pixel size, the more faithfully will be the digital representation of the original analog information. Figure 6.3 shows the results of different pixel sizes on an image. The pixels in Fig. 6.3a are 0.05 mm in height by 0.05 mm in width. The pixels in Fig. 6.3b are 0.11 mm high \times 0.11 mm wide. With the larger pixel size in Fig. 6.3b, some of the fine detail of the image is lost. This is even clearer in Fig. 6.3c where the pixels are 0.22 mm \times 0.22 mm.

¹In microscopy, the Nyquist-Shannon sampling theorem (Shannon 1998) tells us that we cannot accurately detect microscopic information smaller than 2 pixels by 2 pixels of the detector. This is related to the variability in signal collection and not specifically digital imaging, so we will ignore this signal processing issue until the discussion of detectors in Chap. 7.




Fig. 6.4 Pixels per inch and image magnification

Finally, in Fig. 6.3d, the pixels are large enough (0.9 mm \times 0.9 mm) to be visible. Notice that each pixel in Fig. 6.3d contains a single uniform gray tone.²

The number of pixels per inch (or per mm) is a useful value for comparing the amount of information contained in two images. However, in microscopy this term is deceptive. It is only an accurate comparison if the two images are magnified to the same extent. In Fig. 6.4, both images a and b have 400 pixels per inch (PPI). However, if we magnify Fig. 6.4b to the equivalent size as Fig. 6.4a (the result is displayed as Fig. 6.4c), the final rendering only has 67 PPI. Fine details present in Fig. 6.4a are not present in either Figs. 6.4b or c. Since, in microscopy, we are always dealing with magnification of images, pixels per inch is only a useful comparator of the amount of information contained in two images when the microscopy magnification of both images is equivalent.

²Most analog to digital conversions include some smoothing and blending functions, but we will ignore these for the purposes of this discussion.

More small structures than large structures can be placed in a defined area (i.e., they can occur with more frequency). For this reason, we refer to fine detail as high-frequency information and grosser detail as low-frequency information when describing the object being imaged. Under this definition, there is more high-frequency information contained in Fig. 6.4a than in Fig. 6.4b or c.

At first blush, it would seem that you would want to use as small a pixel as possible for imaging. However, the smaller the pixel size, the more pixels are contained in the image. This requires more computer memory for storage and more bandwidth for transferring the image information to the rendering hardware (printer, monitor). In addition, it requires more time to capture the information during the microscope session. All of these parameters need to be considered when determining the optimum pixel size for an application. It is inefficient to store an image at a resolution greater than the capability of the image capture method. The excess empty resolution costs both time and memory. For most applications, you should always attempt to match the resolution requirements with the final use of the image.

We will deal with the capture of images in Chap. 7. Here we will consider practical considerations in determining image storage and rendering. Computer memory has become extremely cheap, so cost is rarely a consideration any longer. However, bandwidth still remains a concern. Thus, it is useful to match the digital resolution to the output requirements. Most modern computer monitors have a resolution in the range of 60–130 PPI (most being between 90 and 100 PPI), although recently some monitors with >300 PPI resolution have come on the market. Thus, if you are only going to display an image on a monitor and the monitor being used is only capable of 100 PPI resolution, it is wasteful to store images at greater than this resolution. The pixels per inch displayed by computer projectors currently use very similar standards to computer monitors, and so the same principles apply. In microscopy, however, monitor display is usually not the only or most important endpoint for our images. We usually will also print our images, enlarge selected areas of our image, and often want to collect quantitative information from our captured images. All of these functions are capable of displaying or analyzing much higher resolution information.

Printers are calibrated in dots per inch. This is not the same as pixels per inch, and the two systems should not be confused. However, for digital printing, a stored image with a resolution of between 300 and 600 PPI is preferred. At this resolution, the human eye will generally not resolve the individual pixels in the print when the print is viewed at a normal viewing distance. Of course, if one gets very close to the image, the individual pixels can begin to be apparent. It is estimated that the human sight can resolve about 1,200 PPI at a very close viewing distance. For normal viewing distances, though, a good rule of thumb is to prepare images at 600 PPI for printing. This is also true of images for submission to journals unless specifically instructed otherwise by the journal's editorial policy.

A computer, of course, does not have the same resolution constraints as the human eye. It sees the image as a series of numbers and so can discriminate differences as small as one unit. So, if the final use of the image will be to have the computer quantify physical characteristics, such as size, area, volume, or co-localization statistics, then one should store images at the maximum resolution feasible taking into

consideration only the sensitivity of the image capture device and any constraints on the resolution imparted by the capture device. Resolution limits imparted by the microscope and image collection method are discussed in Chap. 7.

6.4 Pixel Density

For the analog to digital (A to D) conversion in Fig. 6.2, each pixel was filled with the average of the gray value contained within the pixel area. However, we did not define the range of gray levels that were possible. In computerized digital imaging, the range of possible values (gray levels) is based on binary mathematics, since computers are binary instruments. In binary space, the only values are 0 and 1. In other words, a one digit number in binary space can have a value of 0 or 1. This means that we can only depict two gray levels. If we want to depict a range from black to white, then that means only black or white is possible. By convention, black is usually designated as the lowest value and white as the highest. This situation is depicted in Fig. 6.5a. If, however, we divide our scale from black to white using a two-digit binary number, we have four possibilities: 00, 01, 10, or 11. Each digit is called a bit; Fig. 6.5b depicts the possibilities for a two-bit image. Increasing to two bits expands our range of possibilities so that we can now depict black, dark gray, light gray, and white. Three digits expand the range (bit depth) even further, so we can now represent even more gray levels (Fig. 6.5c). The range increases by 2 to the power of the number of bits. Thus:

- 1 digit = $2^1 = 2$ (1 bit of information)
- 2 digits = $2^2 = 4$ (2 bits of information)
- 3 digits = $2^3 = 8$ (3 bits of information)
- 4 digits = $2^4 = 16$ (4 bits of information)




Fig. 6.5 Gray-scale values for various digital bit levels

As the number of possible values available to define the range from black to white increases, the range encompassed by each single quantum decreases. In fluorescence microscopy, smaller gray-scale quanta allow for the depiction of subtler differences in fluorescence intensity. An 8-bit image allows the depiction of 256 (2^8) gray levels. Based on the ability of the human eye to discriminate differences in contrast, this is considered the minimum required to accurately depict visual information. Of course, machines can discriminate more accurately, so, for quantitative work, even more gray levels are preferable. Many confocal systems are capable of accurately collecting 12-bit images (2^{12} or 4096 gray tones) or even higher bit depths. The advantages of 12-bit images will be discussed in subsequent chapters.

Most microscopists are not facile at working with binary numbers. For this reason, the binary representation of a pixel value is usually converted to a base ten number (Fig. 6.6 a). In the case of a 2-bit image (four gray levels), 0 is used to indicate black, and 3 indicates white. The values of 1 and 2 indicate gray levels based on dividing the range from black to white into four representative gray levels.

In modern computing, 8 bits is termed one byte. Thus, each pixel in an 8-bit images contains one byte of information. If you have an 8 inch by 10 inch gray tone image stored at 600 PPI, the image size would be 28,800,000 pixels. This is because the image will be 4,800 pixels wide \times 6,000 pixels high. This equals 28,800,000 pixels (28.8 megapixels). Since each pixel contains one byte of information, it takes 28,800,000 bytes (28.8 megabytes) to store the gray-scale information of the image.

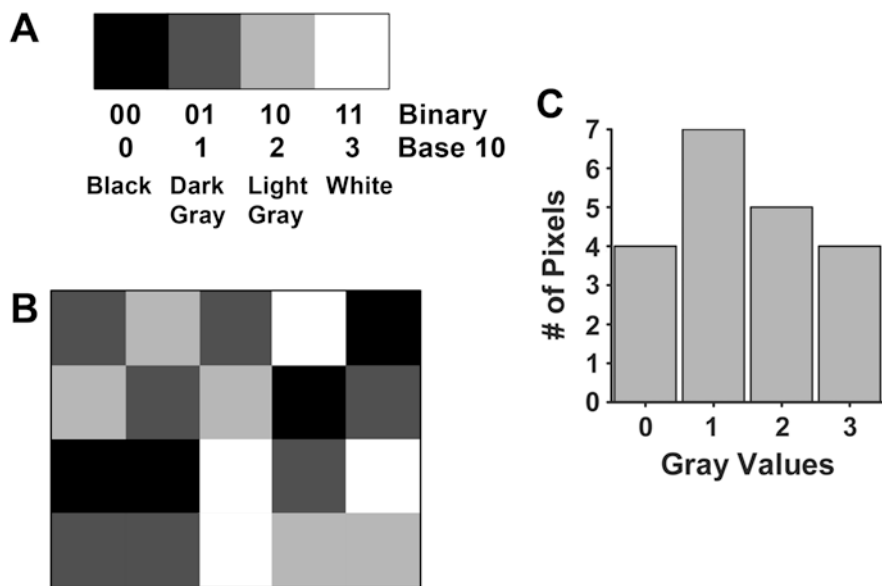


Fig. 6.6 Digital image histograms

6.5 Pixel Histogram Analysis

A useful way to analyze the gray levels in an image is by plotting a histogram of the number of pixels that have a specific gray level (Fig. 6.6b, c). Figure 6.6b shows a two-dimensional array of pixels from a two-bit image (four possible gray values). Figure 6.6c shows a histogram indicating the number of pixels in Fig. 6.6b which have values of 0, 1, 2, or 3. Note that the histogram does not provide any spatial information. We do not know the location of the four black pixels indicated by the bar height for the X value of 0 (black); we have only recorded that there are four black pixels somewhere in the image.

6.5.1 Histogram Stretch

Despite its lack of spatial information, the histogram can be very useful in analyzing and improving an image. Figure 6.7a shows an 8-bit (256 gray levels) image and the histogram from that image. The image has limited contrast with all of the gray level values bunched around the center (medium gray) values. Increasing the difference in density between two adjacent pixels (contrast) can be useful for increasing the visibility of structures. This can be done, based on the gray level (intensity) histogram, by stretching out the values to cover the full dynamic range of 256 gray levels (0 black to 255 white). Figure 6.7b illustrates the increased image contrast that results from spreading the gray values out to use the full dynamic range and also displays the new histogram after spreading.

Most imaging software has at least one routine for reassigning pixel values to stretch or compress the gray level histogram. An important point about histogram stretching for scientific imaging is that histogram stretch does not alter the relative relationship order of one pixel value to another; it only increases the difference in gray tone (contrast) between two pixel values and thus makes the difference more visible to the human eye. Thus, for simply discriminating the location of some structure or examining fine detail, the data contained in the image has not been altered in any significant way; individual structures have simply been made more visible. However, if you are using the absolute pixel gray value as some measure of concentration of one analyte compared to another, you have altered this data by performing a histogram stretch or any other method of reassigning pixel values. For example, in Fig. 6.7b, pixels that originally had a pixel value of 200 now are brighter with a pixel density value 227. Thus, even though procedures such as histogram stretch are relatively benign, they still need to be done with full knowledge of what values are being changed. In looking at the histogram of the stretched image in Fig. 6.7b, you will note that there are now gaps in the histogram. So, although structural information stored in the image file have been made more visible, there are structural details that were not captured. This information cannot be regained. Thus, it is preferable to capture the original image using the full available range of gray values rather than having to extensively stretch the image after collection. Figure 6.7c




Fig. 6.7 Histogram stretch

shows the additional detail that is present when the image is correctly captured. Chapter 9 discusses how to set confocal imaging parameters to insure that the full possible dynamic range is captured in the image.

6.5.1.1 Digital Image Gamma

Gamma is another useful, and relatively benign, parameter that can be used to increase the visibility of specific parts of an image. Gamma is the relationship between the detected gray level (input value) and the gray level that is rendered in the final image (output value) as displayed on the computer screen or in a digital

print (Fig. 6.8). Gamma is a power function. In an 8-bit image, gamma determines the output (rendered value) based on the equation:

$$y = 255 \left(\frac{x}{255} \right)^\gamma \quad (6.1)$$

where y = output value
 x = input value
 γ = gamma




Fig. 6.8 Alteration of image gamma

With a gamma of 1, there is a one-to-one relationship of detected level and rendered level (Fig. 6.8a). An input gray value of 50 would be displayed as 50, and a gray value of 215 would be displayed as 215. In contrast, with a gamma of 0.5 (Fig. 6.8b), an input value of 50 would be displayed as 113, and 215 would be displayed as 234. A gamma less than one helps spread out the information in the darker regions to make it more visible but reduces the difference in intensity between values in the lighter regions. Conversely, a gamma greater than one (Fig. 6.8c) spreads out values in the lighter regions and compresses the darker values. As with histogram stretching, a change in gamma does not change the relative order relationship of the original data. Moreover, with gamma, the original image can be regained by doing the inverse operation³. The inverse is called a gamma correction. Gamma correction routines in some software packages are also called Gamma, so don't be perplexed by this confusing usage. Importantly, changes in gamma do not affect the lowest and highest values. A value of 0 remains 0, and a value of 255 remains 255. Thus, changing the gamma does not alter the dynamic range of the image. When necessary, histogram stretch should be done before changing the gamma.

6.5.2 *Avoid Digital Contrast and Brightness Functions*

Histogram stretch and gamma allow fine control of what detail is displayed in the output image without loss of dynamic range. This is not true of most functions that allow the adjustment of contrast and brightness. These can significantly compress the dynamic range, and we suggest that use of these contrast and brightness functions be avoided. Figure 6.9 demonstrates the effect on dynamic range of a standard method of digital brightness (Fig. 6.9b) and contrast (Fig. 6.9c) control found in a number of popular imaging software packages. To increase contrast, the software shifts the entire register of pixel values toward brighter pixels. However, since you cannot get brighter than white (digital value 255), all the top shifted values become 255. Thus, there is no contrast difference, and what was a difference in the very light gray values in Fig. 6.9a have become white and are lost. To decrease brightness, all the values are reduced. In this case all of the very dark gray pixels become black, and information is lost. Similarly, to alter contrast, the slope of the gamma line is increased or decreased (Fig. 6.9c). Both increasing or decreasing the slope to increase or decrease contrast in the output image results in a loss of dynamic range. To increase contrast, the dynamic range is compressed. To decrease contrast in the output image, values at the low and high ends of the range are converted to black or white, respectively. Once these contrast and brightness functions are run on the image, they cannot be reversed and the lost information cannot be regained. The loss of dynamic range and irreversibility of the functions is why we encourage use of histogram stretch and gamma correction functions to enhance specific details in an image.

³Actually, due to rounding errors, there may be a small difference between the original image and the gamma-correction recovered image, but these differences are generally very small.

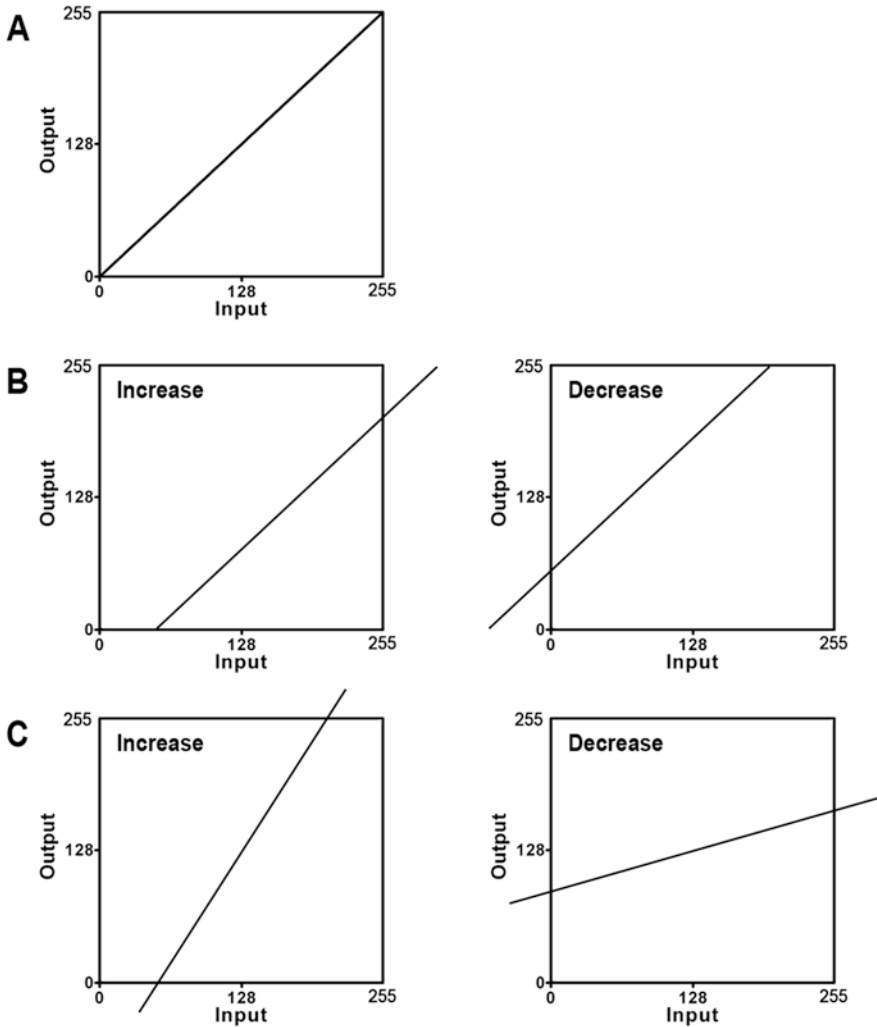


Fig. 6.9 Image brightness and contrast control. (a) Dynamic range of unchanged image. (b) Effect on dynamic range of increasing or decreasing brightness. (c) Effect on dynamic range of increasing or decreasing contrast

6.6 Use of Histogram Stretch and Gamma Changes in Confocal Imaging


Histogram stretch and alterations in gamma are useful tools for making subtle details in an image more visible to a human observer. Gamma changes are particularly useful. The human eye does not have a linear response to light. Our eyes are more sensitive to variations in dark gray levels than equivalent differences in the

brighter registers. The eye responds with a gamma of about 2.2. Scientific digital cameras are generally linear capture devices. This is critical for quantitative work but not necessarily for presenting images to an audience. The makers of computer monitors, printers, consumer-grade cameras, and other digital hardware recognize this and often add a gamma function which converts the linear captured information into a presentation that more closely matches what is seen through the microscope. This is a useful tool, but one needs to be aware of the degree to which a correction is employed when using image information to make quantitative or semiquantitative assessments. The gamma correction employed by a piece of software or hardware is information that is generally not made readily available to the end user. For recreational imaging, it is probably not very important. However, for scientific evaluations, it is always worth digging into the product literature or contacting the manufacturer to determine what corrections a software program, monitor, or printer are imparting to the image data and how to turn those features off, if necessary. Alternatively a test sample with known characteristics can be imaged, and the resultant image examined to determine how the collection software has “corrected” the image. This is a tedious process, but it only has to be done once for each camera system.

In practice, histogram stretch and gamma functions should be used to enhance the scientific usefulness of an image by bringing out additional detail. However, any quantitative analysis or comparison of pixel brightness must be done prior to applying any histogram stretch or gamma functions. Moreover, these and all other image processing steps should be done on non-compressed copies of the original image. The original image should always be stored, unaltered, on archival quality media such as a “write once” CD or DVD. The image should be stored in a lossless image format such as the Tagged Image File Format (TIFF). Image formats are discussed later in Sect. 6.9.

6.7 Image Voxels

A major strength of confocal microscopy is the ability to image defined planes within the sample. Although the planes are thinner than those of a widefield imaging system, they do have a finite height. Thus, confocal images produce 3-D data. The depth of the image (not to be confused with pixel gray-scale depth), combined with the X and Y dimensions, defines a digital volume element (voxel). A voxel is the 3-D analog of the 2-D pixel (Fig. 6.10). In the same way that a pixel represents the average luminance within an area, a voxel’s density will be the average of the luminance values within a defined volume. As such, a voxel represents the smallest unit of volume information that can be depicted in a 3-D digital reconstruction of an analog object. Also, like pixels, the dimensions of voxels are defined in whole multiples of a unit value (quantum). However, again like pixels, the X, Y, and Z quanta defining the voxel do not have to be the same. In most cases, we want to use the

Fig. 6.10 Voxel and pixel relationship

same value for X and Y . However, as discussed in Chap. 7, in microscopy, the Z quanta is usually restricted to a larger value by the physics of microscopy imaging.

Since voxels have three dimensions, voxel information can be used to reconstruct a single plane of the sample, and it also provides information for 3-D reconstructions of larger volumes. If sequential planes are collected, they can be stacked on top of each other to produce a 3-D map of the entire specimen volume. Moreover, the digitized value for a voxel can be transformed in the same way pixel values can be transformed, and the voxel histogram and gamma can be manipulated using the same techniques used for pixels. The methods by which the confocal microscope collects voxel information are discussed in Chap. 9. The constraints on voxel size imposed by the physics of microscope imaging are described in Chap. 7, and the methods of working with voxels to produce digital 3D reconstructions are covered in Chap. 10.

6.8 Color Images

So far we only have dealt with gray-scale images, those that are composed of pixels or voxels with various shades of gray. Digital color microscope images are produced by combining color information encoded as multiple gray-scale images. The two most important methods (color spaces) for microscopists are the Red, Green, and Blue system (RGB color space) and Cyan, Magenta, Yellow, and Key (black) system (CMYK color space). RGB is the mechanism used for computer screens and the one closest to human and microscopic visualization. It is the method implemented on most commercial confocal microscopes. In contrast, CMYK is the system used by most printers.

The RGB system is an additive color system. Black is the absence of any color information, while adding equal contributions of red, green, and blue produces white. Having only red information, and no green or blue, produces red. Likewise, having only green or only blue produces green or blue, respectively. However, by

adding certain proportions of red, green, and blue, we can produce a variety of colors. This is equivalent to combining three beams of different colored light. In practice, in confocal microscopy, three gray-scale images representing red, green, and blue, respectively, are combined. The hue and brightness of each resulting color pixel is the result of the addition of the information contained in each of the three overlain pixels. Thus, for a gray scale of 0–255, an RGB image will have 3 bytes of information for each pixel: one byte for red, one for green, and one for blue. As discussed in Chaps. 7 and 9, in confocal microscopy, we assign the gray-scale density value based on the number of photons collected by the detector, and the hue is assigned arbitrarily but usually based on the specific color filters used to collect the signal.

The CMYK system is a subtractive process. It is equivalent to mixing paint. Conceptually, the process begins with a blank white sheet on which the image will be produced. In the case of a digital image, this means the starting pixel is white. To this white pixel, a certain amount of cyan, magenta, and yellow color is added. Equal contributions of each produce black. Since it mimics the printing process of applying ink to paper, the CMYK system is used by printers. In most printing processes, black ink is substituted for equal amounts of cyan, magenta, and yellow. This saves money on ink and produces deeper dark tones.

Microscopy uses additive light and so is best done in the RGB color space. However, some journals require CMYK images for publication because this system matches the printing process. Although there are similarities between the RGB and CMYK systems, they are not identical. When converting between the two systems, such as when preparing micrographs for publication, the results of the conversion should be carefully checked to make sure the converted image retains accurate color information. Printers also convert the RGB information in a stored image to CMYK when printing. This sometimes requires some readjustment of the images to make sure the printed image is faithful to what was viewed in the microscope or on the computer screen.

A 24-bit RGB color image (8 bits of red information, 8 bits of green information, and 8 bits of blue information) is capable of coding for approximately 16.7 million different colors. Unfortunately, limitations of the capture and rendering devices (cameras, scanners, printers, monitors, etc.) generally are not capable of accurately capturing or depicting 16.7 million different values. Since the deficiencies are not consistent across capture and rendering platforms, this further complicates the accurate reproduction of color information. For most scientific imaging, however, it is sufficient to realize that accurate color transfer is not achievable without a great deal of effort but that this level of accuracy is generally not required for confocal studies, especially since confocal imaging usually collects each set of color information separately. This is explained further in Chaps. 7, 8, and 9.

In order to reasonably accurately depict the colors of the object under investigation, the microscopist needs to understand a few additional concepts about color rendition and color management. These concepts include Hue, Saturation, Brightness, and Luminance.

Hue is what defines the color. All blues have the same hue.

Saturation determines how intense the color appears. A fully saturated color is deep and brilliant. Less-saturated colors appear faded compared to the fully saturated color.

Brightness is a measure of how light or dark the color appears. A brightness value of 0 indicates black.

Luminance describes the perceived brightness. The difference between luminance and brightness is that luminance takes into consideration the color sensitivity of the human eye. For instance, the human eye is much more sensitive to light in the yellow-green range than it is to blue.

Of course, any color in a confocal image is usually artificially added through the use of look-up tables (LUTs). The high-resolution detectors used in modern confocal microscopy detect only the presence of photons. They do not measure the wavelength of those photons and so cannot assign color to the photon. The color component is added to each pixel or voxel by the use of LUTs available in the software based on parameters the microscopist has defined. If a red filter is in place, it is useful, but not obligatory, to display the value for the photons collected as the red information in an RGB image.

6.9 File Formats

The file format is the set of instructions for storing and retrieving the image information. Formats can be broadly categorized into generic and proprietary formats. Proprietary formats are those that are system or software dependent. The Adobe PSD format is an example of a proprietary format. Many of the confocal microscope manufacturer's software also include proprietary formats. The Zeiss LSM and newer CZI format, the Leica LIF, and Nikon ND2 formats are examples. These machine-specific formats allow the software to acquire and store additional information about the image as part of the stored image. In the case of confocal microscopes, this often includes a wealth of information about the microscope settings used to collect the image. Having information about the pinhole diameter, magnification, laser settings, filters, etc. used during image collection is extremely useful. However, the downside of proprietary formats is that the proprietary software is required to decode the stored image. This makes it difficult to share the images with others or to import the images into other software programs. Most confocal microscope manufacturers now also make available a free browser that will recognize and display their proprietary formatted images. Although these free browsers have limited capabilities, they do at least allow basic processing of data and sharing of images with others.

The alternative to proprietary image formats is generic formats. These are formats that are independent of machine, operating system, file system, compiler, or processor. Moreover, images stored in generic formats can be outputted to any hardware

(printers, monitors, etc.) that recognize the format. Most modern image formats also maintain backward compatibility as the format evolves. Luckily, most confocal software also allows saving a copy of an image in one of these generic formats.

There are numerous generic formats, but the two most important for scientific imaging are the TIFF and JPEG (Joint Photographic Experts Group) formats. These formats are recognized by almost all imaging programs, and most proprietary software allows the option of saving in these formats. For scientific imaging, the key difference between the two is that, in the TIFF format, the image information is not changed when the stored image is created. In contrast, the JPEG image format compresses the image so that less computer memory is required to store the image. When computer memory was expensive, compression was very useful. JPEG allows compression of the image at ratios $>1:200$. This compression, though, comes at the expense of image information. The goal of JPEG storage is not to store every bit of information in the digital image. Its goal is merely to store and reproduce the format in a “photorealistic” form. This means that the JPEG image, when rendered from storage, will “look” almost identical to the original. However, since your image is your data, just having an image look like the original is usually not sufficient. Figure 6.11a shows a TIFF image, and Fig. 6.11b shows that same image after conversion to JPEG format. Figure 6.11c shows the subtraction of the two images from each other. All the nonblack pixels in Fig. 6.11c indicate pixels that were altered by the JPEG compression. It should be clear from Fig. 6.11c that scientific imaging should primarily employ lossless storage techniques like TIFF. We suggest TIFF because it is the most serviceable and widely accepted format for scientific image storage and use.

Because it degrades image information, the JPEG format should be avoided for most scientific imaging applications. Although the JPEG format allows one to choose the degree of compression (the higher the compression ratio, the more information that is discarded when storing the image), there is no option in the JPEG format for a lossless saving of the image data. It is also important to note that each time a JPEG image is saved, the program again runs the compression algorithm. Thus repeated saving of an image in JPEG format leads to additive degradation of the image information. The image information loss is the reason why the Ninth Commandment of the Confocal Ten Commandments described in Chap. 1 is “The JPEG image format is Evil.” However, sometimes you need to dance with the devil. This is admissible as long as you recognize the consequences and confirm that the changes occurring do not affect the conclusions drawn from the image. For instance, emailing a 10-megabyte image to a colleague for review may not be accommodated by many email systems. In this situation, JPEG compression to a 50-kilobyte image would be extremely useful. However, quantitative analysis of image information should never be done on a JPEG image. Moreover, as discussed above, it is a good rule of thumb never to resave a JPEG image. When you need to revisit the image, always go back and make a copy of the original TIFF image, and do any post-collection analysis on the TIFF copy.




Fig. 6.11 Effect of JPEG compression on image information. (a) Tiff image of fluorescently labeled pancreas. (b) The same image as in Fig. a but saved in JPEG format. (c) Subtraction of image b from image a showing pixels altered by JPEG compression algorithm

It is worth reiterating here that “your image is your data!” Most modifications done to your image, including JPEG conversion, are irreversible. For this reason, it is critical that you maintain the originally collected image on archival media and only work from copies of the original for any image analysis or conversion routines that alter, in any way, the image information.

Literature Cited

Shannon C (1998) Communication in the presence of noise. Proc IEEE 86:447–457