25

Image Enhancement by Deconvolution

Mark B. Cannell, Angus McMorland, and Christian Soeller

INTRODUCTION

In this chapter we will try to provide a more intuitive (and less mathematical) insight into image formation and practical image restoration by deconvolution methods. The mathematics of image formation and deconvolution microscopy have been described in greater detail elsewhere (see Chapters 11, 22, 23, and 24), so we will limit our discussion to fundamental issues and gloss over most of the mathematics of image restoration. We will also focus on practical ways of assessing microscope performance and getting the best possible data before applying more sophisticated image processing methods than are usually seen in the literature. Before we lose the interest of the confocalist, we would point out there is no such thing as a real widefield microscope because real microscopes have limited field of view and other apertures within the optical train. These apertures introduce a limited degree of confocality *even into widefield microscopes* that, as we will see, improves the behavior of the microscope in a way that makes quantitative image restoration both possible and very worthwhile. The methods (and software) that improve images from conventional widefield (WF) microscopes will also significantly improve confocal images and our practical experience shows that the output of both confocal and multi-photon microscope systems benefit from application of image restoration methods. For those who wish to know the conclusion of this chapter without further reading, it is simply: deconvolve *all* data (if you can). To stress this point, we suggest that even a properly adjusted confocal microscope will not give the best possible confocal image and the application of appropriate deconvolution methods will increase contrast, reduce noise, and even improve resolution. That such worthwhile effects are made possible by computation resides in the fact that during deconvolution the information content of the image can be increased by the addition of information about the imaging process itself as well as that about the statistics of photon capture. Thus, in our opinion, the application of deconvolution methods can always increase image quality, regardless of the source of the image.

BACKGROUND

The resolution of a measurement instrument is always limited so that the recorded *image* of the sample is completely dependent on the properties of the measurement device. Microscopy is simply a measurement process used to quantify information about small objects, and the limits to resolution in microscopy are manifest as blurring in the acquired images. Blurring arises from the intensity of signal at any point in the image being some weighted sum of intensities from all points in the object space. With detailed understanding of the weighting function, we can develop ways of correcting the image for the limiting behavior of the microscope. In addition, almost all microscopic samples are inherently three-dimensional while the detectors [such as the eye, charge-coupled device (CCD) camera, or photographic film] are two-dimensional. Because it is well known that two-dimensional representations of three-dimensional objects lead to artifacts and ambiguity, one should expect fundamental imaging problems in microscopy.

A number of phenomena contribute to the blurring process in light microscopy. First, the wave nature of light dictates that as it passes through apertures (e.g., the optical components making up the microscope), diffraction occurs, which results in spreading and merging of light rays in the image (Abbe, 1873; see Chapter 22). Second, samples can scatter light thereby removing the correspondence between the source of a light ray and its final location in the image plane. This problem makes high-quality imaging in turbid, thick samples (such as brain slices) highly problematic. We should also note that this scattering also prevents certain light rays from passing through the optical train and thereby prevents some information from the object reaching the image. Similarly, a dense object can obscure an (in-focus) object behind it and prevent information from it reaching the image. Finally, refractive index mismatches in the light path (e.g., between the immersion medium and sample) and imperfections in the optical components (and their alignment) can add to aberrations that will also cause light rays to be misdirected.

The conventional WF microscope is designed to correctly image a sample at the focal plane of the objective; however any light source within the conical volume of light collected by the objective lens can contribute light to the image plane. While this is not blurring, it will introduce a loss of contrast for in-plane objects and contribute to the perception of a blurred image. Objects outside the plane of sharp focus are also seen in the image plane as blurred images that also reduce contrast for in-focus objects. (It is the reduction of this effect by confocal microscopes that makes them so useful.) Finally, the limited aperture of the microscope system (i.e., its limited light and information gathering ability) and the limited wavelength of light itself lead to blurring of objects. With so many limiting factors, it is not surprising that real microscope imaging always represents a trade-off between various problems. For example, increasing the numerical aperture of the system may decrease blurring and increase signal from in-focus objects but can also lead to a loss of image contrast due to the acceptance of more scattered light from thick samples. In addition, high

Mark B. Cannell, Angus McMorland, and Christian Soeller • Department of Physiology, University of Auckland, 85 Park Road, Grafton, Auckland, New Zealand

 Handbook of Biological Confocal Microscopy, Third Edition, edited by James B. Pawley, Springer Science+Business Media, LLC, New York, 2006.

numerical aperture objectives have limited working distances that may limit their ability for deep sectioning.

Image Formation

The transformation of information from a real object to blurred image can be expressed mathematically using an operation called convolution. Deconvolution is the reverse operation, whose purpose in microscopy is to remove the contribution of out-of-focus objects from the image plane as well as (partially) reverse aberrations arising from imperfections in the optical train. It should be intuitive that, by using information about the imaging process, deconvolution techniques should be able to improve the quality of the image above that which could be achieved by any other method which does not provide extra information (beyond that contained in the image plane itself). The transformation carried out by the microscope on data from the object can be defined by a mathematical function called the point spread function (PSF). The PSF is the multi-dimensional image of a point in space and it can be practically measured by imaging very small objects (which must be less than the wavelength of light in size) or computed from the physical and optical properties of the imaging system (see Fig. 25.1). Given imperfections in the imaging system (or put another way, limited knowledge of the real optical system) the latter approach is inherently limited, but assumptions about the properties of the PSF can help some blind deconvolution methods.

Put simply:

$$i(x, y, z, t) = o(x, y, z, t) \otimes psf(x, y, z, t) \quad \textbf{(1)}$$

where x, y, z, t are the dimensions of space and time and i is the recorded image, o the actual underlying object. The $\otimes$ operator is convolution (see below and Fig. 25.2). If we take the Fourier transform ($F\{\}$) of this equation, the $\otimes$ is replaced by multiplication:

$$F\{i(x, y, z, t)\} = F\{o(x, y, z, t)\} \times F\{psf(x, y, z, t)\} \quad \textbf{(2)}$$

This simple equation shows that image restoration should be possible and might be achievable by dividing the Fourier transform of the image by the Fourier transform of the PSF and then taking the inverse Fourier transform. Before exploring this idea further, it should be noted that these equations remind us of the importance of correct data sampling (image and PSF) and the behavior of the PSF over time and space. For example, although we generally assume that the microscopes do not change properties as we focus through the sample and from day to day, is this really so? When objectives are damaged by lack of care (such as from a collision with slide/stage or accumulation of dirt), the gradual loss of imaging performance may go undetected. This is particularly a problem for confocal microscopes whose pinhole can effectively hide a poor or deteriorating PSF at the expense of signal strength. Also note that coverslips (which are an integral part of the optical system) vary in thickness from batch to batch and therefore introduce variable amounts of spherical aberration.

From here on, we assume that everything is time invariant. (This is not a trivial assumption as, in reality, samples are *never* time invariant, they move — especially if alive, they bleach, the microscope focus drifts with temperature, etc.) So, from Eq. 2, if one can measure or compute the PSF with sufficient accuracy, why can images not be restored to the point where any desired resolution is achieved? The major problem resides in the noise that is always present in a physical measurement. The convolution or blurring of the object by the PSF is spatial frequency–dependent and attenuates high-frequency components (provided by small features and edges) more than low frequency (large, smooth) objects. Hence, deconvolution must boost high spatial frequency components more than low spatial frequency components. Because noise

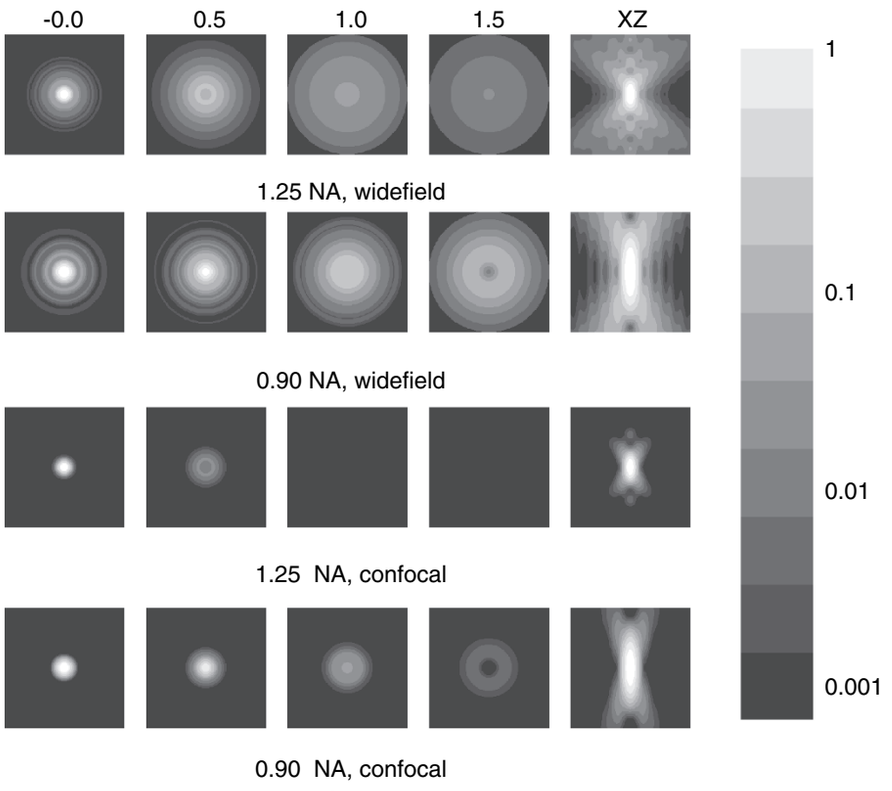

FIGURE 25.1. Theoretical PSFs calculated using widefield equations. On the left, sections at 0.5 µm z-intervals from the plane of sharp focus are shown while xz sections are shown on the right. Panel (A) shows PSFs for two typical lenses. Note the effect of reducing NA on the axial extent of the PSF. Panel (B) shows equivalent PSFs for the above lenses in a confocal configuration. Note that the confocal PSF looks very similar to the widefield PSF with the outlying wings removed. Again, the improved axial response of the higher NA lens is quite clear in these images. The gray scale is shown on the right and is logarithmic.

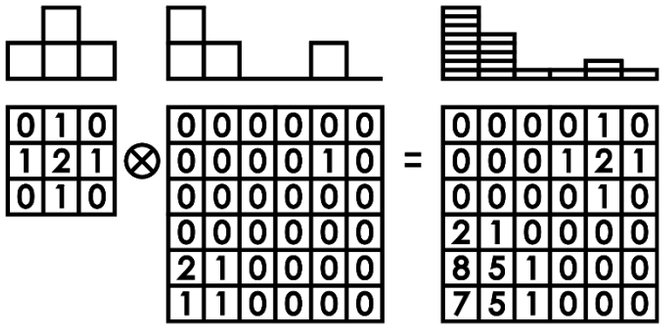

FIGURE 25.2. Schematic diagram demonstrating the convolution (⊗) operation with a 6 × 6 pixel object and a 3 × 3 pixel blurring kernel. The profiles above show the maximum projection of the two-dimensional grids as would be seen looking across the planes from above. Note how the contrast of the peaks in the image is reduced and smeared across the image.

is present at all frequencies in the image right up to the spatial sampling frequency which is 1/(pixel size), noise components are boosted at all frequencies during deconvolution. At more than half the sampling frequency of the image, there is no real object data retained (from the Nyquist sampling theorem) but the ever-present noise receives an even larger boost during deconvolution.

Put another way, because noise does not come directly from the object, but is introduced during the imaging process, and because noise is rich in high-frequency components, deconvolution (which seeks to boost high-frequency content) can amplify noise to the point where it masks, and renders useless, any information in the resulting deconvolved image. The addition of noise also generates ambiguity in the image restoration (deconvolution) process so that more than one optimal restoration solution exists with no *a priori* method of determining which solution best represents the real object. Current deconvolution algorithms work around some of these limitations, often by making a number of reasonable assumptions about the object (such as smoothness and nonnegativity) and include extra information about the noise process itself.

Having introduced image restoration by deconvolution, we will now describe the convolution and deconvolution processes in more detail and give a guide to some of the practical issues regarding image deconvolution. We will also discuss how some popular deconvolution algorithms vary, and assess the utility of image restoration. We will show that deconvolution can substantially improve WF microscopy, to the point where it can be used as a viable and practical alternative to confocal laser scanning microscopy (CLSM) and can also be used to good effect to improve the quality of confocal and multi-photon microscope images.

FORWARDS: CONVOLUTION AND THE IMAGING SYSTEM

Before it is possible to develop an algorithm for reversing the defects induced during the imaging process, it is necessary to develop a basis for image formation. This section will give a shortened version of that formalization and more detailed descriptions can be found in Agard and colleagues (1989) and Young (1989). The notation used generally follows the conventions of Press and colleagues (1992).

The starting point for this formulation requires two assumptions regarding the imaging process: linearity and shift invariance. The principle of linearity is met if the sum of the images of two separate objects is the same as the image of the combined objects. The assumption of linearity is generally safe in fluorescence microscopy provided that detectors are linear and we avoid self-quenching and self-absorption by the fluorophores (Chapters 16 and 17, *this volume*). Shift invariance implies that the image of an object will be the same regardless of where in the field of view that object lies. (While no real imaging system meets the requirement of shift invariance, it is a reasonable assumption for a high-quality research microscope, subject to our being aware of the potential complication introduced by changing aberrations across the field and with focus.) Now, because any object can be represented by a superposition of many delta functions (point light sources) whose individual images are simply shifted copies of the PSF, it follows from the principles of linearity and shift invariance that the whole image is made from the superposition (sum) of appropriately scaled and shifted instances of the PSF. For clarity, we can express this in one dimension as an equation:

$$i(x) = \int_{-\infty}^{+\infty} o(x - x')psf(x')dx' \qquad \textbf{(3)}$$

which is, in fact, mathematically the same as the convolution of the object with the PSF. In symbolic form:

$$i(x) = o(x) \otimes psf(x) \qquad \textbf{(4)}$$

which gives a basis for the first general equation given above. As noted above, convolution in the spatial domain is equivalent to multiplication in the frequency domain so we can write (for three-dimensional objects):

$$F\{i(x, y, z)\} = F\{o(x, y, z)\} \times OTF(\varsigma, \psi, \zeta) \qquad \textbf{(5)}$$

where OTF() is the optical transfer function and is the Fourier transform of the PSF (see Figs. 25.3 and 25.4), ξ, ψ, and ζ are the spatial frequency coordinates of the OTF in Fourier space, derived from the PSF sampled in x, y, and z, respectively. The OTF is introduced here because convolutions are most quickly computed in the frequency domain by using the fast Fourier transform (FFT). Because the PSF is assumed to be shift invariant, only a single Fourier-transformed PSF is used for computations and the values of the PSF in the spatial domain *per se* have no practical value for deconvolution. Thus, it is more convenient to describe the PSF of a microscope in terms of its OTF rather than its PSF. The OTF directly describes how blurring affects the various frequency components making up an object and immediately gives insight into the resolving power of the microscope. Although we usually measure the PSF in the image plane (rather than at the rear aperture of the objective where the OTF exists) and most researchers are more familiar with an ideal PSF and can identify microscope defects by inspecting it (see below), the OTF is actually more useful. Put another way, because the performance of the entire microscope is the result of the combination of the OTFs of each optical lens and aperture, it is easier to immediately determine the effect of an aperture or lens in a microscope from its OTF. As an example, because a circular aperture reduces the amplitude of the OTF with increasing spatial frequency, it is immediately clear that the ultimate determinant of resolution must be the most limiting aperture in the system [usually the numerical aperture (NA) of the objective lens].

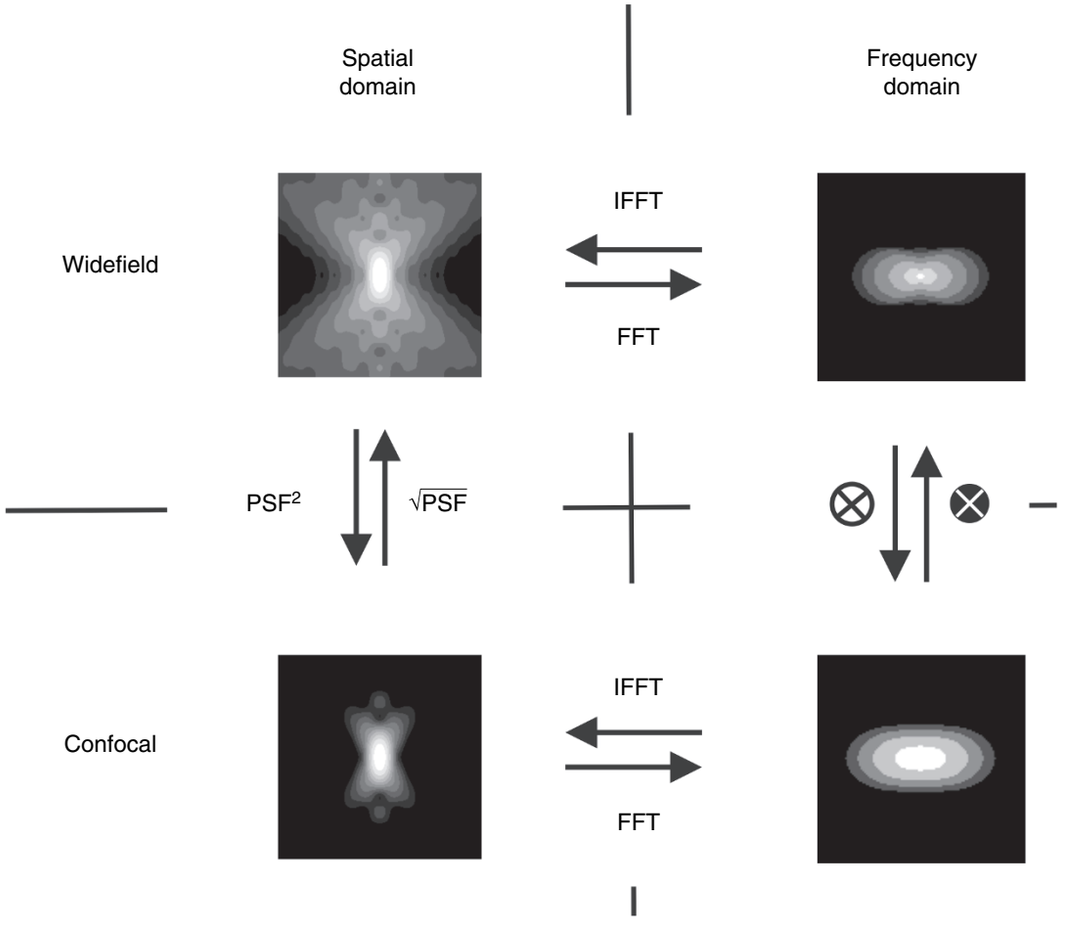

FIGURE 25.3. The relationship between the PSF and OTF in confocal and widefield modes of a microscope. For simplicity, we assume that the pinhole function leads to a simple squaring of the PSF, a reasonable approximation for an optimal pinhole with similar excitation and emission wavelengths. The convolution of the widefield OTF with itself in the frequency (Fourier) domain is the equivalent of squaring in the spatial domain. Note that the inversion of the convolution (inverted convolution symbol) is difficult in the frequency domain and is more easily accomplished in the spatial domain (by taking the square root). This illustrates that some operations are more simply carried out by operating in the appropriate domain; convolution is easier in the frequency domain while multiplication is simpler in the spatial domain. Such interchangeability between spatial and frequency domains is central to deconvolution methods.

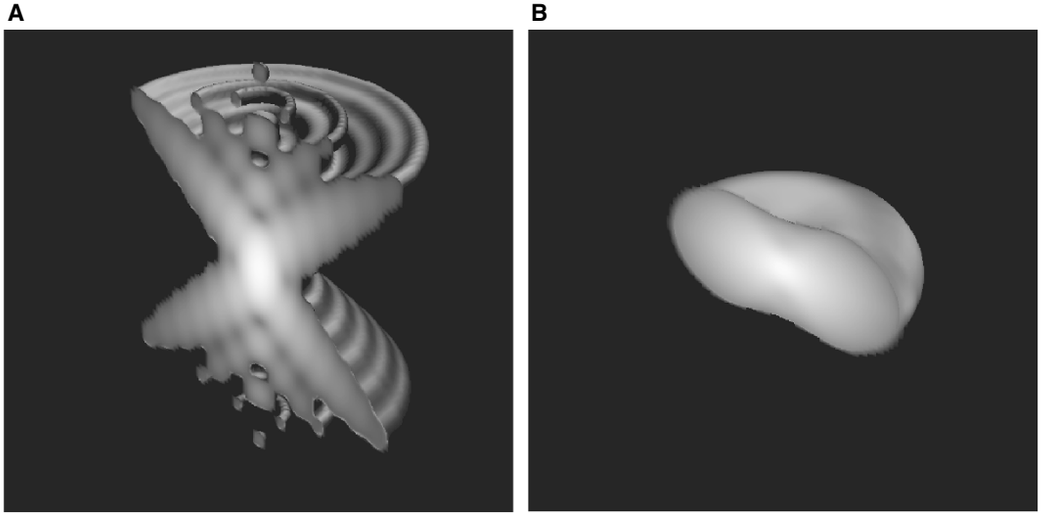

FIGURE 25.4. Volume rendered illustrations of the widefield PSF (A) and OTF (B). The PSF extends to infinity as a cone with ripples in the intensity in all directions. The OTF appears smaller being closer to a toroidal shape except that the center values are non-zero. This leads to a cone of low values in the center of the toroid. The toroid width in both directions is limited by diffraction. The gray scale shading is proportional to the logarithm of the intensity.

PROPERTIES OF THE POINT SPREAD FUNCTION

Because the PSF of a microscope determines how it creates images of an object, the shape of the PSF determines the way in which images of objects blur into each other in the final image. At the focal plane of a WF microscope, the PSF [see Fig. 25.1(A)] exists as a central Airy disk, surrounded by dimmer concentric rings that are the result of diffraction occurring as light passes through the circular apertures of the microscope. Away from the plane-of-focus, the PSF spreads outwards, forming two apex-opposed cones with diffraction ripples within them [Fig. 25.4(A)]. This pattern reflects the fact that WF microscopes form images from cones of light gathered by an objective lens, with the angle of the cone and thus the shape of the PSF being determined primarily by the NA of the objective lens (Chapters 11, 22, 23, and 24, *this volume*). The axial extent of the concentric cones of the widefield PSF at z-levels well removed from the focal plane show how significant the out-of-focus contribution is to a WF image. This is not the case for the CLSM PSF whose PSF approaches background levels with distance [Fig. 25.1(B)], a point which will be revisited later.

It is important to note that, unlike theoretically derived widefield PSFs, measured widefield PSFs are rarely symmetric around the focal plane, nor radially symmetrical about the optical axis. Asymmetry along the optical axis (z-axis) is commonly due to spherical aberration, which may result from refractive index mismatches between the objective, immersion medium, and sample or tube length/coverslip thickness errors. Radial asymmetries in the PSF are commonly the result of misalignment of optical components about the z-axis, either as tilt or decentration (see Cagnet *et al*., 1962; Keller, 1995) (see Fig. 25.5).

The extent of the PSF determines the resolution of the microscope. PSF size is often expressed in terms of the width at which it is half the maximum intensity (full-width half-maximum, or FWHM). A typical FWHM for the widefield PSF of a high power immersion objective lens might be <0.4 µm in the focal plane. The axial extent of a true widefield PSF should be infinite with constant energy at all planes, but real microscopes lose signal energy with distance because additional apertures (beyond the entrance pupil) in the system lead to the more-distant wings of the widefield PSF being clipped (or vignetted). Nevertheless, the extended axial response of the WF microscope results in a massive loss of in-plane contrast for extended specimens which are either fluorescent or scattering.

QUANTIFYING THE POINT SPREAD FUNCTION

Most deconvolution methods (i.e., not blind ones) require knowledge of the PSF relevant to the particular imaging conditions encountered. PSFs can be estimated either mathematically or from direct measurement. In blind deconvolution methods an estimate of the PSF is also produced by the deconvolution algorithm. (Note that if the object and image are known, the OTF is given by the Fourier transform of the image divided by the Fourier transform of the object; see Eq. 5.) While it is relatively easy to quantify the PSF of a microscope under controlled conditions (see Fig. 25.5), matching the conditions encountered during an experimental imaging situation is often impractical, and commonly some approximation to the likely real PSF is used.

Most methods for calculating theoretical PSFs are based on equations in the definitive work of Born and Wolf (1980). A good description of the calculations for a confocal PSF is given by van der Voort and Brakenhoff (1990), in which the PSF is calculated

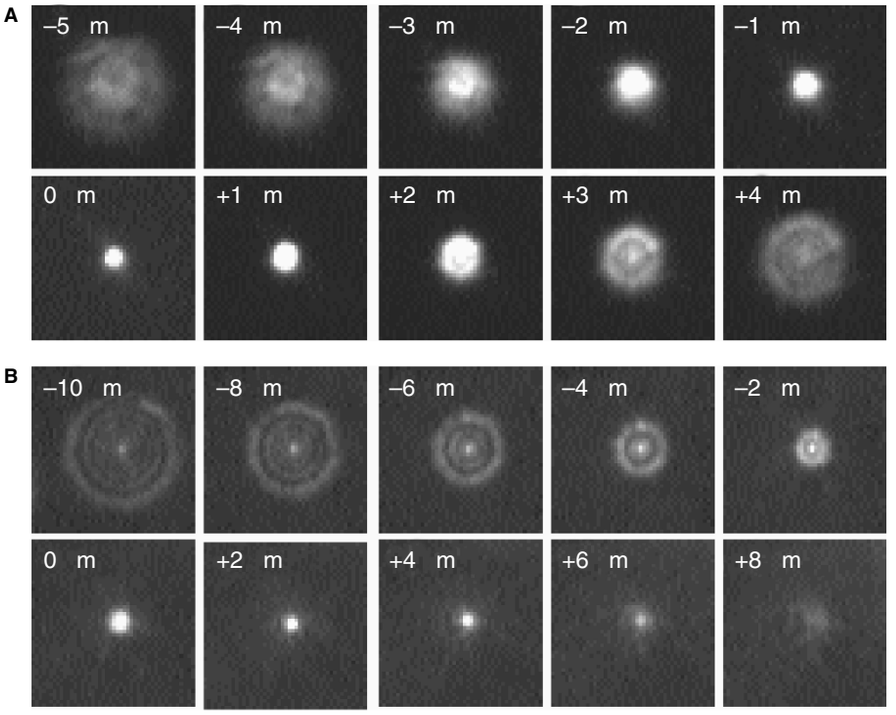

FIGURE 25.5. Measuring the widefield PSF. (A) shows the appearance of a 0.2 µm bead with varying degrees of defocus. Note that although the most focused image is reasonably circular, with defocus extra structure appears in the Airy disk. These uncontrolled aberrations need to be measured for correct deconvolution. Note also that the aberrations are not symmetrical about the point of sharp focus so that enforcement of symmetry in blind deconvolution would be problematic. (B) shows a through-focus series with spherical aberration. Asymmetry above and below best focus is evident.

from the NA of the objective, the illuminating and emitted wavelengths, and the refractive index of the immersion medium in either (simpler) paraxial forms or with WF integrals (see Fig. 25.6). Numerically derived expressions for the PSF give an indication of the best possible resolution for a given objective but these limits are not achievable. The real world limitations in alignment of the optical path and lens aberrations inevitably lead to a real PSF that is larger and has more asymmetric structure than predicted by mathematical approaches. Particular attention should be paid to whether a paraxial approximation is appropriate because there are considerable differences between the results of paraxial and WF calculations for the high-NA objectives which are often used in experiments (Fig. 25.6). The problem is confounded by software packages that do not reveal what method is used to calculate the PSF. In our experience, real PSFs are typically >20% bigger than calculated versions (based on the stated NA) and their shapes are rarely perfectly symmetrical about the optical axis and focal plane, even when we have tried to control refractive index mismatches. It should be intuitively obvious (from the preceding discussion) that the accuracy of the deconvolved image can be no better than that of the PSF that was used for the deconvolution. Errors in the PSF can produce bizarre artifacts in the image (such as hollow cones or slanting tails around bright objects) as well as the algorithm not properly converging to minimum noise (see Fig. 25.7).

The PSF of a microscope can be measured by taking images of commercially available fluorescent beads of <0.2 µm diameter. These, being smaller than the diffraction limit of visible light, act like point light sources, and their imaged size and shape therefore represent the microscope PSF. It is a common misconception that even smaller beads give better estimates of the PSF. This is not the case as the imaging process is a convolution and the convolution is dominated by the larger of the two objects being convolved. In addition, the beads become much less bright as smaller beads are used (note that the signal intensity will fall with the cube of the diameter) so precise measurement of the entire PSF becomes very difficult, especially as the PSF is recorded away from the focal plane. For those who feel uncomfortable using 0.2 µm spheres, it is possible to use deconvolution of the PSF image using a 0.2 µm sphere (the actual object) to yield the precise PSF (from Eq. 5), but in practice it is rarely necessary to go to such lengths. In our experience, the difference in the final results of deconvolution using bead images and PSFs derived by deconvolution of bead images with spherical sources is negligible.

When imaging beads for the purposes of PSF measurement, the image data should be obtained at as high a resolution as possible. In this case, Nyquist's criterion of using twice the desired resolution is insufficient, and pixel resolutions 0.05 µm (or even better) are desirable. This has some advantages: (1) The fine data spacing facilitates data resampling to enable the PSF resolution to match that of the recorded image. It is often the case, especially when confocal variable zoom controls are used, that the final image resolution may be different from that used to acquire the PSF. (2) Generous over-sampling also allows a greater dynamic range to be recorded as the received photons are recorded in more voxels. Even if the detector is only 8 bit, over-sampling by a factor of 2 in all three dimensions raises the effective bit resolution to 11 bits. This increase in effective dynamic range is very useful as it enables both the bright center of the PSF as well as more dim distant features to be recorded. For deconvolution, low noise PSFs are required and this can be achieved by averaging a number of individual bead images, possibly from different beads (because the principle of shift invariance requires that the bead image — the PSF — is not affected by the position of the bead). This approach requires an algorithm for aligning the centers of the image of each bead so that the averaging operation can combine data points in each bead image. The coordinates of the center of the 3D bead image can be derived by calculating the center of signal mass of an image that has been subjected to a thresholding procedure to completely remove any background signal. The most distal parts of the image should also be removed by the chosen threshold to prevent the estimate of the bead centroid depending on the position of the bead within the sampled data volume.

Perhaps the most problematic aspect of measuring PSFs using fluorescent beads is matching the experimental imaging conditions to which the PSF is to be applied. First, to account for the effect of any chromatic aberrations (as well as the direct effect of wavelength on resolution), PSFs should be measured from beads that emit light of a similar wavelength to the experimental fluorophore being imaged. (Beads are available with a range of emission wavelengths from 415 to 680 nm.) Errors in refractive index, either between the objective and immersion medium or between immersion and bathing/mounting media, need to be consistent for correct replication of the aberrated PSF. Because experimental samples are often less well controlled than calibration bead samples, we often add a few beads to critical experimental samples to provide a check of the actual PSF in the biological sample. These beads are readily recognized in images as they usually stick to the outside of cells (in live-cell imaging) and do not bleach as rapidly as the experimental sample. Although it is possible to remove the bead from the experimental image digitally, we have never had to do this in practice as it is usually possible to find a suitable field of view free of beads.

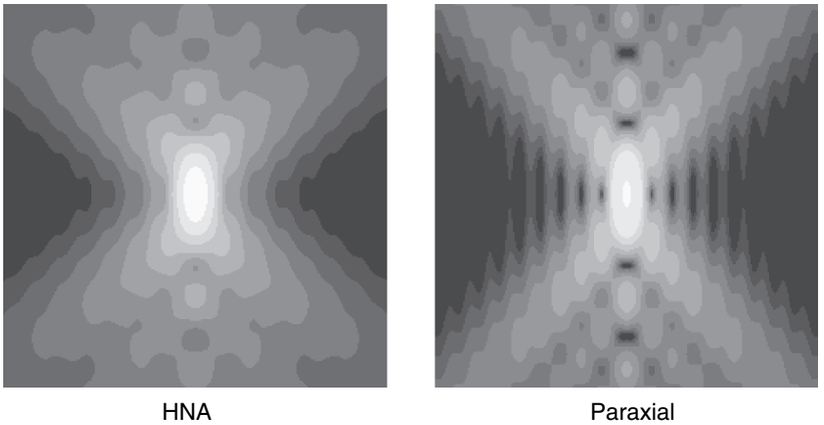

FIGURE 25.6. Comparison of widefield and paraxial computations of the PSF. The left panel shows the computed PSF for a 1.25 NA objective using the integrated field equations (van der Voort and Brakenhoff, 1990). The right panel shows the equivalent results for the paraxial equations which combines the electric vectors using the assumption that $\sin(\theta) = \theta$. The paraxial assumption leads to greater ringing in the response but a narrower Airy disk.

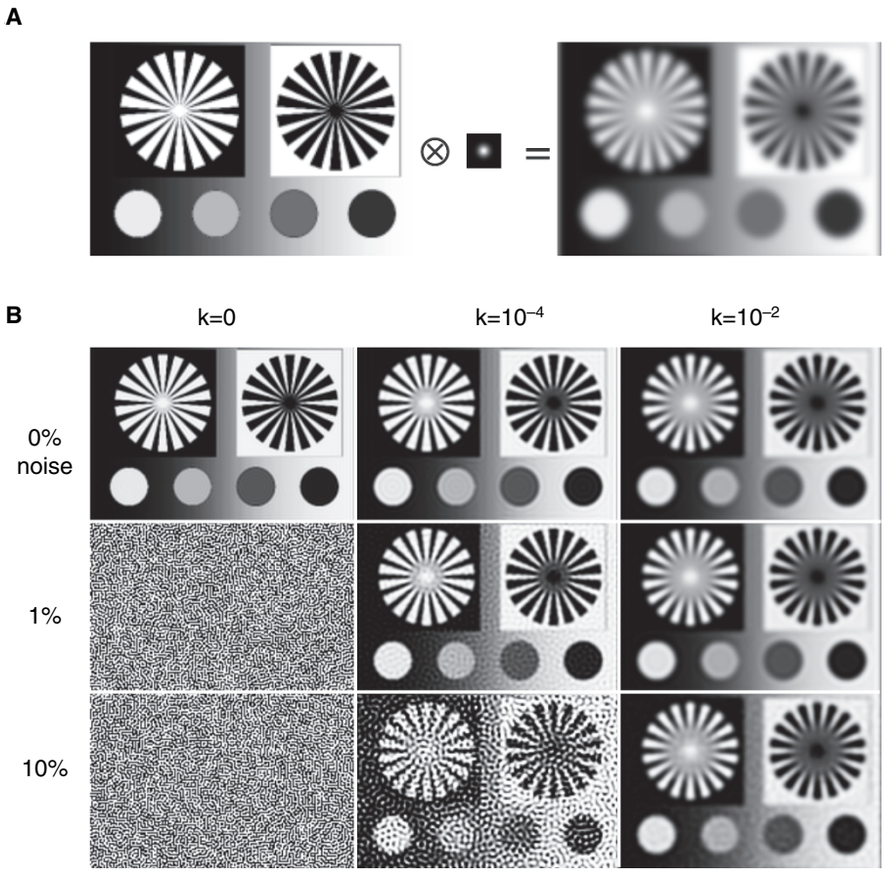

FIGURE 25.7. Image restoration by Wiener filtering. Panel (A) shows the original test image and the PSF used to blur it. The resulting blurred image has lost contrast for all small features. Panel (B) shows the results of image restoration with variable amounts of noise added to the blurred image for various values of k (see Eq. 10). On the left side, without noise the inverse operation perfectly restores the data. However, even with 1% noise ($N = 0.01$) inverse filtering yields fewer data than even the original blurred image. When $k = 10^{-4}$, the effect of noise is reduced but at the expense of high frequency performance. The loss of high frequency components leads to ringing near sharp edges and can be seen in the small colored circles. When $k = 5 \cdot 10^{-2}$, noise is even better controlled but the loss of high frequency information is more severe. This can be clearly seen by comparing the centers of the spoked targets. Note that the image where $k = 10^{-4}$ and $N = 0.01$ is comparable in appearance to the image $k = 5 \cdot 10^{-2}$ and $N = 0.1$. Thus increasing k controls increasing noise at the expense of reduced spatial resolution.

THE MISSING CONE PROBLEM

As noted above, while the PSF provides information about the spatial throughput of a microscope, its Fourier transform (the OTF) gives information about the frequency response of a microscope directly. While the NA of the microscope limits the highest spatial frequency that can be recorded, extended sources (i.e., lower spatial frequency objects) lose signal strength with defocus because source light rays must originate in the finite field of view of the microscope. Thus, the OTF of the WF microscope looks like a toroid [Fig. 25.3 and 25.4(B)] with so-called missing cones above and below the focal plane. These missing cones of reduced signal reflect the cone of light acceptance of the objective (assuming it is the limiting aperture in the microscope). The width of the torus in all planes is limited by the numerical aperture of the objective giving rise to a bandlimit equal to $2NA/\lambda$ (as required from Nyquist sampling of the object using the spatial frequency of light in the immersion medium). The axial band limit is $NA^2/2\eta\lambda$ while the angle of the missing cone is $\arcsin(NA/\eta)$, where η is the limiting refractive index of the mounting media between objective and specimen.

In a CLSM, the pinhole aperture restricts the ability of light from out-of-focus objects to reach the detector. This means the light-focusing pattern of the objective is used to determine both the illumination and emission light paths, and so the PSF of a CLSM is approximately the square of the widefield PSF (assuming that the ratio of excitation to emission wavelength is close to 1). The PSF of a two-photon microscope is similarly related to the power of the widefield PSF because the probability of near simultaneous multiple photon absorptions is proportional to the illuminating intensity raised to the nth power, where n is the nearest integer ratio of illumination to excitation wavelengths (Denk *et al.*, 1990). As a result, the PSFs [Fig. 25.2(A–C)] are ovoids with the long axis corresponding to the optical axis of the microscope. The corresponding OTF does not have missing cones and is closer to an ovoid (Fig. 25.3). This can be explained by the fact that the confocal PSF is the product of the illumination and detection PSFs (Wilson, 1990) and in Fourier space, multiplication becomes convolution. Convolution of one toroid by another tends to fill in the center void.

As pointed out earlier, deconvolution can be thought of as the division of the Fourier transform of the image by the OTF of the microscope. The low values of the OTF in the neighborhood of the missing cone and near the bandlimit greatly amplify any noise in the image and this severely limits the ability of deconvolution methods to restore images. This problem is partly ameliorated by the apertures in the optical train that limit the ability of the microscope to capture light from distant objects (whose in-plane components would have low OTF values); their attenuation prevents the massive amount of signal (and noise) in the integrated far field from completely swamping in-focus data. Put another way, blurring of data from nearby objects is the biggest problem, while signal from more distant objects is attenuated because real microscopes are not true WF instruments. This problem is reduced in the CLSM because the missing cone does not exist, however, the loss of photons due to the pinhole reduces the signal-

to-noise ratio in the image and, during deconvolution, all noise is amplified.

NOISE

Photon noise is always present. The quantal nature of light means that a given illumination intensity is associated with statistical variation in the time at which photons are detected. The number of photons observed over any period behaves as a Poisson process whose variance is equal to the mean. In fluorescence microscopy, the number of photons that can be collected is limited by photodamage (bleaching) and the signal-to-noise ratios are quite low. In addition, although the PSF and object are generally considered to be three-dimensional, the bleaching processes, fluctuations in illumination intensity, and even movement of the object make real objects four-dimensional. To remove this fourth dimension we need to take a three-dimensional snapshot to freeze time-dependent changes. This implies limited exposure periods must be used, which of course further limits light collection.

Photons need to be converted to electrical signals for quantification and this process also introduces noise. Even in the absence of photons, all detectors give spurious signals (dark current) while undetected photons lead to a reduced signal-to-noise ratio. Over a range of wavelengths, average detector efficiencies are generally below 50% and this figure is further degraded by significant losses in the optical components in front of the detector (such as barrier filters and mirrors). It is not unusual to see overall detection efficiencies well below 5% in real microscope systems at some wavelengths and in the face of such signal losses the photon (shot) noise in the system can be severe. In addition, the detector may add electrical noise to the signal before digitization.

A third source of noise resides in the digitization of the signal. In most systems, the detector provides an analog signal that is digitized into discrete levels, and 8- to 12-bit detectors are most common. There is, of course, a one-half–bit uncertainty in the digitization of any such signal and this typically leads to between 0.2% to 0.01% noise. Although this may seem small (compared to photon noise), this noise is also amplified during data processing. (To appreciate the potential contribution of this component we should recall that the original in-focus data may represent <5% of the total signal.) The digitization noise does not apply to all digital photon-counting systems. Despite the concerns, photon noise generally remains the dominant noise source (Mullikin *et al*., 1994; Pawley, 1995; van Kempen *et al*., 1997) in microscopic fluorescence images, at least when using PMTs. When using CCD detectors, other sources of noise are added. For a very good CCD with a readout noise of ±10 photoelectrons, readout noise will dominate until 100 photons are detected (at which point these noise sources are equivalent). For cameras, another source of noise is added due to pixel-to-pixel variation in gain (and/or sensitivity). This flat-field noise is age and temperature sensitive, it increases in direct proportion to signal strength, and it might be equivalent to photon noise at a signal strength of ~5000 photons. At high light levels this noise source dominates but can often be partially controlled by correcting the image for the gain/sensitivity variation (by normalizing the image to a reference image of a white field).

Finally, the thousands of computations required during deconvolution may lead to computational errors due to truncation. Fortunately, modern computers have quite long word lengths so this problem is largely disappearing, but errors in programming such as inappropriate typecasting can lead to truncation errors propagating. Computational errors are often hard to diagnose as they usually appear as general noise in the image.

DECONVOLUTION ALGORITHMS

In the previous discussion we have shown that in fluorescence microscopy (in all its modes including widefield, confocal, and multi-photon) the imaging process can be mathematically described by a convolution. Convolution consists of replacing each point in the original object with its blurred image (the PSF) in all three dimensions and summing together overlapping contributions from adjacent points to generate the resulting three-dimensional image. As pointed out above, the inverse operation or deconvolution can be performed (at least in principle) by dividing the Fourier transform of the recorded image by the Fourier transform of the PSF (this argument followed from Eq. 6). Although mathematically correct, this inverse filter algorithm does not work well in practice because the inevitable noise in the recorded data can be so strongly amplified as to render the deconvolved data useless. It is this problem that more sophisticated deconvolution algorithms have been developed to overcome and more robust methods are well established. In this section we will describe the ideas behind these improved methods without too much mathematics. Note that when using the term image here, we mean the full three-dimensional data set, often called an image stack. If desired, the reader can consult one of the *Numerical Recipes in X* texts (where *X* may be Fortran, C or C++; see also http://www.nr.com for online access) for more detailed discussion of some of the fundamental mathematical concepts involved (i.e., maximum likelihood, minimization of functions, etc.).

Nearest-Neighbor Deconvolution

Deconvolution algorithms vary greatly with respect to computational intensity, that is, how long it actually takes to compute the deconvolved image. Among the simplest algorithms that are of practical use are the nearest-neighbor approaches. Such algorithms are based on the simplifying assumption that the loss of contrast in any given plane is only due to signal arising from objects in planes immediately above and below it. To restore the data with this assumption, we can subtract appropriately blurred versions of the object planes above and below the current focal plane from the recorded image data. The blurring of the adjacent planes is, of course, determined by the PSF. For example, the signal contribution from object structures in the plane immediately above can be obtained by a two-dimensional convolution:

$$o_{j+1} \otimes psf_{j+1} \qquad \textbf{(6)}$$

where, assuming that image planes are recorded a spacing Δz apart, psf_{j+1} denotes the values of the PSF in the next image plane, Δz from the focus. To calculate this contribution we would need to know o_{j+1} (in addition to the three-dimensional PSF), but we only have image data at this point. We therefore make a second approximation (which is no worse than the first) that $o_{j+1} = i_{j+1}$. Thus, a fraction of this blurred nearest-neighbor data is then subtracted from the data recorded in the current image plane, or, as an equation:

$$o_j = [i_j - c(i_{j-1} \otimes psf_{j-1} + i_{j+1} \otimes psf_{j+1})] \qquad \textbf{(7)}$$

where the index j identifies the current plane, $j - 1$ and $j + 1$ the planes immediately below and above, o and i are the deconvolved object and recorded image, respectively, and c is a constant <1.

This algorithm increases in-plane contrast and blurring in the focal plane can also be reduced by applying a 2D inverse filter of the in-focus PSF (psf_j^{-1}). For further details see Agard and colleagues (1989):

$$o_j = [i_j - c(i_{j-1} \otimes psf_{j-1} + i_{j+1} \otimes psf_{j+1})] \otimes psf_j^{-1} \quad \textbf{(8)}$$

Note that all the convolutions are two-dimensional, making this method suitable for real-time implementation with FFTs. On the other hand, one should not be surprised that the enhancement obtained with this method is only moderate and (most important) not quantitative. These limitations reside in the subtractive nature of the process which will add noise and alter signal amplitudes. Nevertheless, where qualitative imaging at high speed is required, useful increases in in-plane contrast can be achieved with this method. For better results one needs to resort to fully three-dimensional algorithms.

An even faster implementation can be made by assuming that the data in adjacent planes is the same as that in the image plane, that is, $i_{j\pm1} = i_j$. In this no-neighbors method, the PSF is used to subtract a blurred contribution of the image plane from itself. Strictly speaking, this is not a deconvolution method at all but is formally equivalent to an unsharp-mask high-pass filter. Clearly, the gross assumption underlying this approach limits its applicability to very sparse three-dimensional objects and any intensity quantitation would be out of the question.

Wiener Filtering

Turning now to full three-dimensional approaches, as noted above the object can be deconvolved by inverse filtering, that is, direct division of the Fourier transform of the data by the OTF:

$$F\{o\} = \frac{F\{i\}}{OTF} \quad \textbf{(9)}$$

In this equation, the low values of the OTF in the region of the missing cone will increase noise in the restored image. Put another way, the data which is lost in the missing cone cannot be replaced unless we supply some other information that acts as a constraint on the restoration. Nevertheless, while the amplification of noise makes direct inversion nearly useless (at least for microscopic images), it can be improved by a simple modification (Shaw and Rawlins, 1991):

$$F\{o\} = F\{i\} \times \frac{OTF}{OTF^2 + K} \quad \textbf{(10)}$$

where K is an additive correction factor (see Figure 25.7). By adding K, noise amplification will be limited where the OTF is small. It can be seen that if $K = 0$, the equation reduces to the original inverse filter (Eq. 9). The impact of K can be reduced in the low frequency domain (where noise is less of a problem) by making it depend on the spectral characteristics of noise and object data. In this case, K can be mathematically optimized and this optimal filter is called the Wiener filter. Unfortunately, Wiener filters do not generally provide optimal image quality because: (1) suppression of any part of OTF during restoration affects the entire image, (2) where the OTF is small, noise will be amplified, possibly to an unacceptable degree, and (3) small errors or uncertainties in the OTF will lead to large errors in the reconstruction. Generally, for microscopic imaging purposes, Wiener and similar linear filters have been superseded by nonlinear iterative deconvolution algorithms that have become more applicable with increasing computer power.

Nonlinear Constrained Iterative Deconvolution Algorithms

Why use nonlinear methods? The problem resides in the fact that linear methods cannot guarantee that the deconvolved object is positive. Rather, in the presence of noisy image data, some pixels in the deconvolved object will assume negative values, a result which is obviously incorrect. For fluorescent images all photon fluxes must be positive so nonlinear algorithms that ensure no negative data values in the reconstructed object are clearly preferable. Such positivity constraints can be easily enforced within iterative algorithms. In iterative algorithms, the deconvolved object is not computed in a single pass (as is the case in nearest-neighbor or Wiener filtering) but the object is calculated in a series of steps that are then repeated. Each step results in a new estimate of the object that should be somewhat closer to the actual object. During this process, positivity can be enforced by explicitly setting negative values in the current estimate to zero. The small error that is introduced by this clipping procedure should be corrected in the next iteration and, with a well-behaved algorithm, the estimate converges to a nonnegative solution that closely approximates the true object. After a certain number of steps, the process is terminated by the user to yield the deconvolved data.

All iterative algorithms start from an initial estimate and we will discuss reasonable starting choices briefly later. The algorithm then enters the iterative loop in which a new estimate is calculated from the current estimate and the recorded image data. The positivity constraint is then enforced by setting all negative pixels in the new estimate to zero. New estimates are calculated repeatedly using this scheme until the user (or the software implementation that he uses) decides to terminate the computation and accept the last calculated estimate as the deconvolved object.

Expressed as a piece of pseudo code:

1. Pick a starting estimate o^0
   a. Calculate OTF (if blind)
2. Compute an initial new estimate (o^{k+1}) from previous estimate o^k, recorded image data i and OTF
3. Set all negative pixels in o^{k+1} to zero to give the current best estimate and apply any other constraints
   a. If blind, estimate a new OTF from F{i}/F{o^{k+1}} applying any constraints such as NA, symmetry and nonnegativity.
4. Go to step 2 unless a stop criterion is reached

Differences between iterative algorithms primarily arise from the way in which the new estimate is calculated (step 2).

To calculate a better estimate of the object, we must have some way to find out if our current object guess (o^k) is good (or not). Ideally, the estimate blurred by the PSF would be the same as our recorded image i except for some small difference due to the noise in the image data. Following this idea, we can calculate an error signal as the difference between our blurred guess and the actual image data. So to update the estimate we can take a fraction of error signal and add it to the current estimate (a Newton iterative approach). This approach underlies the Jansson–van Cittert algorithm (Jansson *et al*., 1976). As an equation:

$$o^{k+1} = o^k + \gamma(i - (o^k \otimes psf)) \quad \textbf{(11)}$$

and with a good choice for γ (which can vary across the image) this scheme should lead to successively better estimates. However, in early implementations the choice of γ was rather *ad hoc* (Jannson, 1970; Agard *et al*., 1989) and noise leads to uncertainty as to whether the algorithm actually achieves the optimal solution.

In more recent iterative algorithms, the updating scheme is derived from more stringent mathematical reasoning. One class of algorithms that work well in practice are called maximum likelihood algorithms. The idea behind this approach is based on the notion of the likelihood of a certain object o to underlie our recorded data i. In this context, we can view the deconvolution problem as a multi-dimensional fitting problem. In this view we anticipate that for the correct choice of fitting parameters (actually the pixels values in o), there should be a finite probability that our data should have been observed. Thus, we have to try to find the object o which *maximizes* the likelihood that our data would have occurred given a specific noise model (which is required to mathematically calculate these probabilities). Put simply, the idea is to fit a function (the blurred object guess) to the data with appropriate weighting for noise (which should be a concept familiar to all experiment scientists); see also the discussion in Press and colleagues (1992, section 15.1).

To deal with noisy fluorescence data, we can therefore calculate the probability that our data would have occurred, assuming the noise follows a Poisson (or some other noise model) distribution. Having found an explicit mathematical expression for this probability we then proceed to find the object o that maximizes this probability. Instead of trying to find this maximum directly we can make life a bit easier by looking for the maximum of the logarithm of the probability. Due to the properties of logarithms, this log-likelihood is maximal at the same location that the likelihood itself is maximal, however, the mathematical expressions generally become easier to handle after taking the logarithm. To actually find the maximum, mathematicians use the observation that at a maximum the derivative of the log-likelihood will be zero. Cutting through the mathematical derivation, one finally arrives at an expression for the updating step in the presence of Poisson noise:

$$o^{k+1} = o^k \times IF\left\{F\left\{\frac{i}{o^k \otimes psf}\right\} \times OTF'\right\} \qquad \textbf{(12)}$$

where $IF\{\}$ is the inverse Fourier transform and OTF' is the complex conjugate of the OTF. (The Fourier transform yields real and imaginary parts so the OTF is complex.) By definition, the complex conjugate of a complex number $x + iy$ is $x - iy$. Multiplication by the complex conjugate in Fourier space is known as correlation. Note that in this algorithm we do not use the additive error $I - o^k \otimes psf$ but rather the relative difference $i/(o^k \otimes psf)$. This algorithm is also known as the Richardson–Lucy algorithm after the authors who first derived this scheme for astronomical applications (Richardson, 1972; Lucy, 1974). For an example of the image enhancement possible with this algorithm see Figures 25.8 and 25.9. Assuming Gaussian noise in the data (for large photon fluxes, the Poisson noise model approaches Gaussian noise) we arrive at an alternative updating step that uses the additive error signal:

$$o^{k+1} = o^k + IF\{F\{i - o^k \otimes psf\} \times OTF'\} \qquad \textbf{(13)}$$

With both types of maximum likelihood algorithms, we should stop after a certain number of iterations for best results as, eventually, the error and/or noise starts to increase again.

This brings us back to the start and stopping choices that must be made with iterative deconvolution algorithms. Reasonable choices for the first estimate are (a) the recorded image itself, (b) a constant estimate with every value set to the mean of the recorded image, and (c) a smoothed version of the recorded image. Using the Richardson–Lucy maximum-likelihood algorithm (Eqs. 12, 13) choices (b) or (c) are recommended. A more detailed discussion can be found in van Kempen and van Vliet (2000). We also need a criterion of when to stop the iterative procedure and accept the current estimate as the deconvolved object. A reasonable approach involves monitoring the relative change between estimates as the iteration proceeds. As the relative difference

$$\frac{o_{k+1} - o_k}{o_k} \qquad \textbf{(14)}$$

between subsequent estimates falls below a chosen threshold (that generally depends on the amount of noise in the images) the iterations should be stopped. For example, with some of the data shown in this chapter, thresholds of $\sim 10^{-4}$ were used.

There are other issues that should be considered in a serious implementation of constrained-iterative deconvolution methods, for example, background estimation (van Kempen and van Vliet, 2000) and prefiltering of data and PSF (van Kempen *et al.*, 1997). The other approach to iterative deconvolution is to use a regularization term, often the Tikhonov functional, as set out by Tikhonov and Arsenin (1977), which aims to limit the generation of large intensity oscillations in the deconvolved image (which will grow as the loss of high frequency information in the OTF becomes limiting). A number of similar methods based on this approach exist, and these vary primarily in the approach taken to find the solution to the regularization term which minimizes the signal oscillations or roughness in the deconvolved image. One such solution is the iterative constrained Tikhonov–Miller (ICTM) algorithm, which uses a minimization strategy based on conjugate gradients to minimize the Tikhonov functional. Tikhonov-regularized forms of the Richardson–Lucy algorithm also exist, representing a consolidation of these two classes of solution.

Comparisons of a number of deconvolution algorithms (van Kempen *et al.*, 1997; van Kempen, 1999) show that there is relatively little difference between their various deconvolved results. The regularized algorithms tend to be less sensitive to noise than Richardson–Lucy-based methods, and reach an appropriate stopping point sooner. However the RL algorithm provides slightly better deconvolution results, as quantified by the mean square error between the image and the PSF-blurred estimate of the deconvolved object. A useful improvement in the Richardson–Lucy method can be achieved by prefiltering the image and PSF with a Gaussian smoothing filter which reduces high frequency noise contributions without affecting the deconvolution of lower frequency image components.

One other variant of the general deconvolution algorithm is blind deconvolution (Holmes and O'Connor, 2000 and Chapter 24, *this volume*), where the deconvolution is performed in the absence of an independently measured (or calculated) PSF. Instead, the PSF is iteratively estimated along with the experimental data during the deconvolution process simultaneously. Proponents of blind deconvolution suggest this method is superior to using PSF-dependent calculations with a PSF rendered incorrect by sample-induced aberrations and scattering. On the other hand, one can argue that by supplying less information to the reconstruction, the final result must be inferior, if the PSF was known to sufficient precision. In addition, if the blind method makes use of symmetry in the PSF and/or the NA of the objective, artifacts will be produced if lens aberrations are present and/or the assumed NA is modified by, for example, improper illumination of the rear aperture of the lens.

Comparison of Methods

A common question is: What method works best? Unfortunately, there is no single answer because all imaging methods represent a

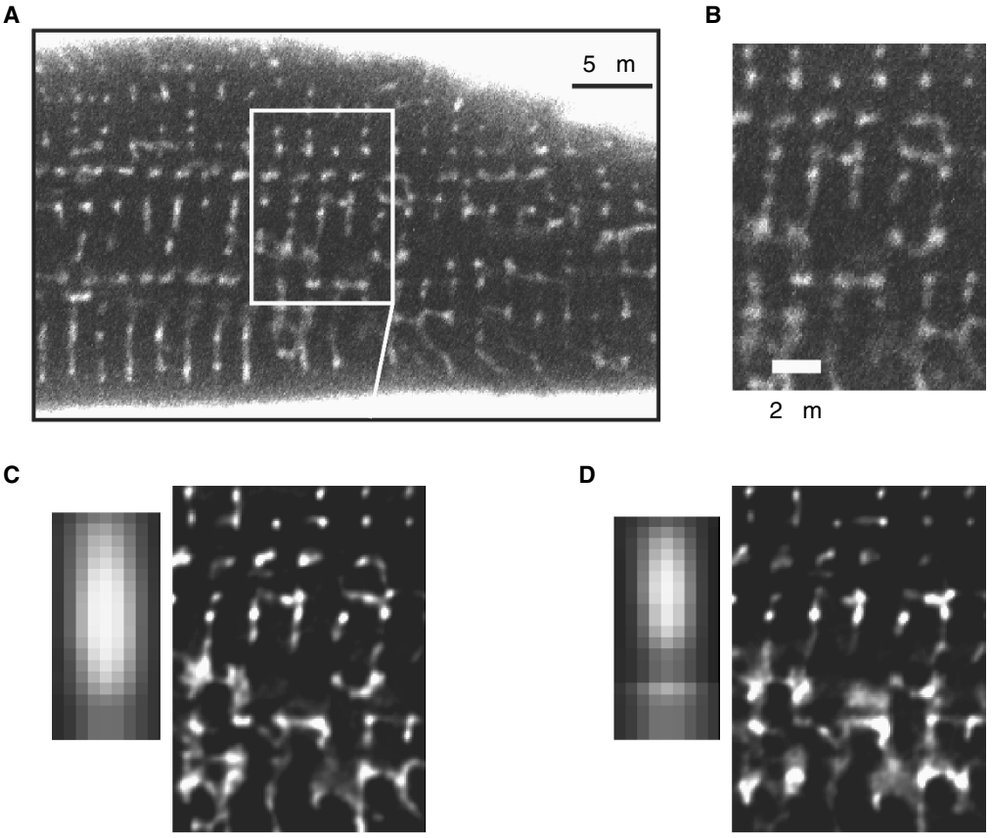

FIGURE 25.8. Maximum likelihood deconvolution of the cardiac *t*-system. (A) shows a two-photon image of a single, rat cardiac cell in dextran-linked fluorescein. The fluorescein has filled the *t*-tubules which are <~300 nm in diameter. (B) shows an enlarged view from a part of (A). Note the relatively low signal-to-noise present in this image although tubules can be clearly seen. (C) shows the restored image. Note the improvement in signal-to-noise ratio as well as the increased contrast. There has also been a useful but moderate increase in spatial resolution. The fine structure of the *t*-system can now be appreciated. (D) shows the effect of deconvolving the image with an incorrect PSF. At first sight, the results appear reasonable, but closer inspection reveals the presence of artifactual data when compared to (C). In view of the general problem of controlling spherical aberration, without good controls such artifact could easily lead to erroneous conclusions.

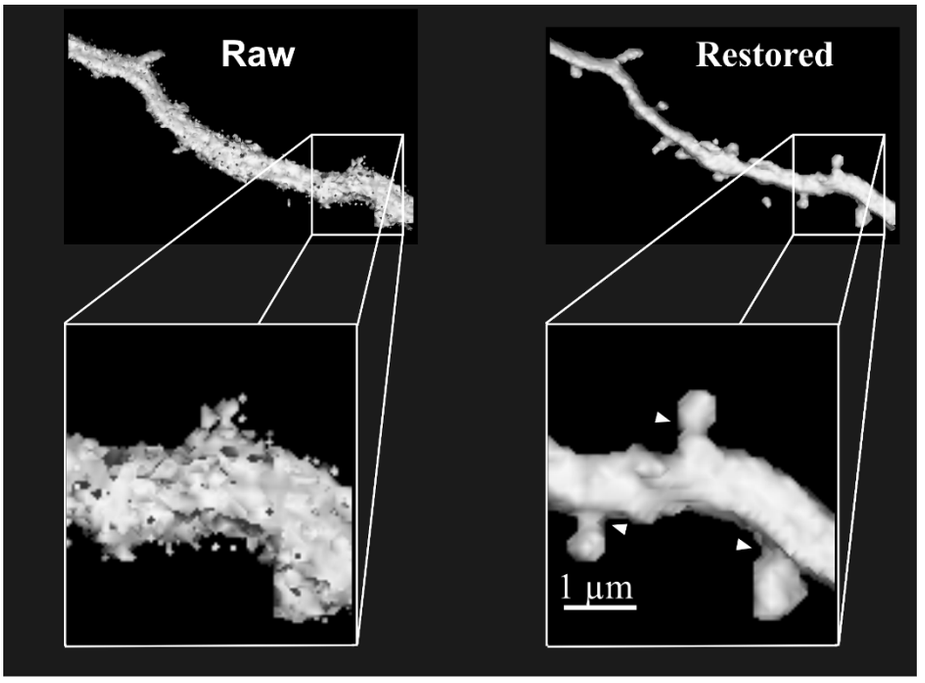

FIGURE 25.9. Enhancement of two-photon images by deconvolution. The panel at the left shows a neuronal dendrite in a living brain slice. The limited photon yield leads to little fine structure being visible in the volume-rendered image. After maximum likelihood deconvolution, synaptic spines are clearly visible and even appear to have necks (*arrows*). In this case, deconvolution has been much more useful to control noise rather than increase optical resolution.

trade-off between the three fundamental limits in information transfer processes: resolution versus speed versus noise. Thus, improvements in one area will usually compromise another and it is up to the experimenter to consider what is actually desired of the imaging method. This principle is generally limiting even within a chosen restoration method, because once resolution has been increased to a certain point, further gains can only be made by compromising noise and/or speed. As a guide, we suggest that deconvolved confocal/multi-photon images provide the ultimate image in terms of resolution (see Figs. 25.8 and 25.9). In some cases, deconvolved WF images with iterative methods can achieve comparable resolution to confocal/multi-photon methods. Nearest-neighbor methods provide a lower level of image improvement that can help control loss of contrast in WF images but suffer from their non-quantitative nature (i.e., one should not compare brightness between different parts of the image). Deconvolution of WF images may provide a resolution approaching that of confocal methods but would be problematic in scattering specimens or in the face of spherical aberration.

In terms of speed, nearest-neighbor methods are fastest with execution times for interative-constrained methods being comparable regardless of whether the source images are WF or confocal/multi-photon. This pecking order also seems to apply in terms of signal-to-noise ratio (including errors in absolute signal levels) with nearest-neighbor methods being worst and iterative constrained methods being best. In terms of noise, WF methods are initially superior until resolution starts to become limiting, at which point differences between methods largely disappear.

Finally, because even after deconvolution, image quality is critically dependent on the quality of the input data, we suggest 20 tips for better imaging:

1. Store unused objectives lens down, clean, and dry in the manufacturer's case.
2. Before using an objective, inspect it with a microscope eyepiece for cleanliness and damage — note any scratches across the metal front of the lens — possibly resulting from collisions with stage or specimens. If spring-loaded, check the objective front is not locked down.
3. Check the coverslips for thickness (check the coverslip requirement on the microscope barrel, usually 0.17 mm) and that they are suitable for the chosen objective. If the objective has a correction collar, use it!
4. Get a supply of subresolution fluorescent beads in a wide range of excitation and emission wavelengths to measure the PSF and test microscope performance. Mount them in standard slides (avoid mountants with harsh solvents); you can also purchase slides already made up.
5. Oversample the PSF and subsequently resample to image resolution by applying a Gaussian filter (to reduce noise) followed by a cubic-spline interpolation. If appropriate, apply a circular-symmetry constraint on the PSF to further reduce noise.
6. Add a small quantity of beads to samples before sealing the coverslip to enable visual inspection of the PSF in widefield.
7. Always seal wet-mounted specimens to prevent them drying out.
8. Minimize spherical aberration by appropriate selection of objective lenses/mounting media. If the objective has a correction collar, adjust it to make the observed PSF as symmetrical as possible. This can be achieved by watching the bead image as the microscope is slightly defocused. The Airy disk should grow symmetrically in both directions with defocus (see Fig. 20.3, *this volume*).
9. If a high-power oil lens must be used with a water-based specimen, try raising the refractive index of the specimen with sucrose, mannitol, or salt if osmotic pressure is not a problem, but be careful not to allow crystal formation. It should be noted that an 84% weight/volume sucrose solution has the same refractive index as glass (~1.5). Alternatively, use immersion oil having a refractive index >1.52.
10. In confocal applications, always open the pinhole slightly to increase light flux while maintaining a reasonable degree of confocality to achieve the desired axial resolution.
11. Use the highest voxel resolution possible in xy and z, but try to avoid bleaching problems.
12. Minimize bleaching problems with antifade mountants and latest-generation fluorochromes.
13. If imaging multiple wavelengths sequentially, image the longest wavelength fluorochromes first.
14. Ascertain the minimum and maximum extents of the specimen at the lowest possible illumination intensity (with the pinhole quite open in confocal applications). The idea is not to produce a nice image at this point but simply detect where the fluorescent specimen resides.
15. Keep a copy of the lens PSF with the image data and record voxel dimensions, objective, mounting media (some imaging software allows comments to be placed in the image header).
16. Image stacks will usually fit within a single CD-ROM (do not use any data compression unless you are sure it is lossless) so make a CD-ROM copy of the data before leaving the laboratory.
17. Avoid mechanical artifacts during imaging by keeping acoustic noise down, the door closed, and avoiding touching any part of the microscope system. Vibration isolation systems may be essential in laboratories with suspended floors. Air-cooled laser system fans should be mechanically decoupled from the microscope with a duct to a separate exhaust fan.
18. Use dimmable incandescent lighting in the laboratory and dim the lights while collecting data. On inverted microscopes, cover the specimen with a black cup.
19. Use a lookup table that makes it easy to spot over- and underflows in the data. Make sure the black level of the imaging system is set to >0; this is even more important than having a few pixels saturated.
20. Always have sample slides containing beads on hand and always check the microscope image quality with them before starting imaging, every day.
21. Deconvolve everything!

REFERENCES

Abbe, E., 1873, Beitrage zur theorie des microscopes und der mikroskopischen wahrnehmung, *Schultzes Arch f. Mikr. Anat*. 9:413:468.

Agard, D.A., Hiraoka, Y., Shaw, P., and Sedat, J.W., 1989, Fluorescence microscopy in three dimensions, *Methods Cell Biol*. 30:353–377.

Born, M., and Wolf, E., 1980, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, Pergamon Press, Oxford, England.

Cagnet, M., Francon, M., and Thrierr, J.C., 1962, *An Atlas of Optical Phenomena*, Springer-Verlag, Berlin.

Denk, W., Strickler, J.H., and Webb, W.W., 1990, Two-photon laser scanning fluorescence microscopy, *Science* 248:73–76.

Holmes, T.J., and O'Connor, N.J., 2000, Blind deconvolution of 3D transmitted light brightfield micrographs, *J. Microsc*. 200:114–127.

Jansson, P.A., Hunt, R.H., and Plyler, E.K., 1970, Resolution enhancement of spectra, *J. Opt. Soc. Am*. 60:596–599.

Keller, H.E., 1995, Objective lenses for confocal microscopy, In: *Handbook of Biological Confocal Microscopy* (J.B. Pawley, ed.), Plenum Press, New York, pp. 111–126 or Chapter 7, *this volume*.

Lucy, L.B., 1974, An iterative technique for the rectification of observed distributions, *Astronomy J*. 79:745–765.

Mullikin, J.C., Van Vliet, L.J., Netten, H., Boddeke, F.R., van der Feltz, G.W., and Young, I.T., 1994, Methods for CCD camera characterization. *Proc. SPIE* 2173:73–84.

Pawley, J.B., 1995, *Handbook of Biological Confocal Microscopy*, 2nd ed., Plenum Press, New York.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., and Flannery, B.P., 2002, Numerical recipes in C: The art of scientific computing. Cambridge University Press, Cambridge, UK, 1–970; available free, on-line at www.library.cornell.edu.

Richardson, W.H., 1972, Bayesian-based iterative method of image restoration, *J. Opt. Soc. Am*. 62:55–59.

Shaw, P.J., and Rawlins, D.J., 1991, Three-dimensional fluorescence microscopy, *Prog. Biophys. Mol. Biol.* 56:187–213.

Tikhonov, A.N., and Arsenin, V.I., 1977, *Solutions of Ill-Posed Problems*, V.H. Winston, Washington, DC.

van der Voort, H.T.M., and Brakenhoff, G.J., 1990, 3-D image formation in high-aperture fluorescence confocal microscopy: A numerical analysis, *J. Microsc*. 158:43–54.

van Kempen, G.M., 1999, Image restoration in fluorescence microscopy, In: *Advanced School for Computing and Imaging*, Delft University of Technology, Delft, The Netherlands, p. 161.

van Kempen, G.M.P., and van Vliet, L.J., 2000, Background estimation in nonlinear image restoration; *J. Opt. Soc. Am. A-Opt. and Im. Sci.* 17(3): 425–433.

van Kempen, G.M., van Vliet, L., Verveer, P., and van der Voort, H., 1997, A quantitative comparison of image restoration methods for confocal microscopy, *J. Microsc*. 185:354–365.

Wilson, T., 1990, Confocal microscopy, In: *Confocal Miscrocopy* (T. Wilson, ed), Academic Press, London, pp. 1–64.

Young, I.T., 1989, Image fidelity: Characterizing the imaging transfer function, *Methods Cell Biol*. 30:1–45.