A Workingperson's Guide to Deconvolution in Light Microscopy

BioTechniques 31:1076-1097 (November 2001)

Wes Wallace[1], Lutz H. Schaefer[2], and Jason R. Swedlow
University of Dundee, Dundee, Scotland, [1]Brown University, Providence, RI, USA, and [2]Advanced Imaging Methodology Consultation, Kitchener, Ontario, Canada

ABSTRACT

The fluorescence microscope is routinely used to study cellular structure in many biomedical research laboratories and is increasingly used as a quantitative assay system for cellular dynamics. One of the major causes of image degradation in the fluorescence microscope is blurring. Deconvolution algorithms use a model of the microscope imaging process to either subtract or reassign out-of-focus blur. A variety of algorithms are now commercially available, each with its own characteristic advantages and disadvantages. In this article, we review the imaging process in the fluorescence microscope and then discuss how the various deconvolution methods work. Finally, we provide a summary of practical tips for using deconvolution and discuss imaging artifacts and how to minimize them.

INTRODUCTION

Deconvolution is a computational technique for improving the contrast and resolution of digital images. It includes a suite of methods that seek to remove or reverse the blurring present in microscopes images caused by the limited aperture of the microscope objective lens. Nearly any image acquired on a digital fluorescence microscope can be deconvolved. In addition, new applications to transmitted light images are now available (29). Three-dimensional images made up of a series of optical sections are particularly well suited for improvement by deconvolution.

Deconvolution is often described as an alternative to confocal microscopy. This is not strictly true since confocal images can themselves be deconvolved. However, most users apply deconvolution to images recorded on a standard "wide-field" fluorescence microscope. This approach yields images of comparable resolution to a confocal microscope (49). In fact, confocal microscopy and wide-field-deconvolution microscopy both work by removing blur, but they do so by opposite means. Confocal microscopy prevents out-of-focus blur from ever being detected, by placing a pinhole between the objective lens and the detector, through which only in-focus light can pass (30). Wide-field microscopy allows blurred light to reach the detector; deconvolution then attempts to either subtract blurred light from the image or reassign it back to its source (1,37,53). Confocal microscopy is especially well suited for thick specimens such as embryos or tissues, while wide-field-deconvolution microscopy has proven to be a powerful method for imaging samples requiring low light levels, such as living cells bearing fluorescently labeled proteins and nucleic acids (9,18,24,35,46,54).

Our goal in this article is to introduce deconvolution to the working biologist at a level that is more practical than theoretical, but more rigorous than a user's manual. Because of space constraints, we focus on the application of deconvolution to 3-D wide-field images of fluorescent biological specimens.

Causes of Image Degradation

Image degradation can be divided into four independent phenomena: noise, scatter, glare, and blur (7,56). The principal task that deconvolution sets for itself is to remove blur. Deconvolution algorithms can and do remove noise, but this is a relatively simple aspect of what they do.

Noise is a quasi-random disarrangement of detail in the image, which in its most severe form has the appearance of "white noise" or "salt-and-pepper noise", the kind of signal degradation seen in broadcast television during bad reception. We call it "quasi-random" because the statistical distribution of noise can be predicted if the mechanics of its source are known. In digital microscopy, the source is either the signal itself (so-called "photon shot noise") or the digital imaging system. The mechanics of both sources are understood; therefore, the statistical distribution of noise is known. Signal-dependent noise is characterized by a Poisson distribution, while imaging system noise usually follows a Gaussian distribution. Thus, noise in the image can be removed by appropriate filters, and most deconvolution software includes "preprocessing" routines that accomplish this. The topic of noise in digital microscopy is discussed elsewhere (31,39–41).

Scatter is a random disturbance of light caused by its passage through regions of heterogeneous refractive index within a specimen. The effect of scatter is a truly "random" disarrangement of the image detail. No completely satisfactory method exists yet to predict scatter in a given specimen. However, we can say that the degree of scattering depends on the thickness of the speci-

men and on the optical properties of the specimen material. The thicker the specimen, the more scatter there is; and the more heterogeneous the refractive index of the specimen material, the more scatter there is.

Glare, like scatter, is a random disturbance of light, but occurring in the lenses or filters of the imaging system rather than within the specimen. The level of glare in the modern microscope is minimized by the use of lenses and filters with antireflective coatings.

Blur is a nonrandom spreading of light caused by its passage through the imaging system and lenses. The cause of blur is diffraction, and an image whose resolution is limited only by blur is considered "diffraction-limited" (7, 31,36). This is an intrinsic limit of any imaging system and is the determining factor in assessing the system's resolution limit. Optical theory includes sophisticated models of blur, and, with modern computer power, we can apply such a model to digital images-this is the basis for deconvolution. Because of its importance in deconvolution, the theoretical model of blur will be introduced in greater detail below. However, we emphasize here that all imaging systems cause blur independently of whatever other forms of image degradation are caused by the specimen or the electronics. It is precisely this independence of blur from other types of degradation that makes it possible to remove blur by deconvolution.

Bohren (6) points out that scatter, glare, and blur have the same physical cause, namely the interaction of light and matter. However, the composition and arrangement of molecules in a given material (whether glass, water, or protein) gives each material its particular optical properties. For our purposes, what distinguish scatter, glare, and blur are the location where they occur and the possibility of generating a mathematical model for them. Because scatter is a local, irregular phenomenon occurring in the specimen, it is difficult to model [although see Kam et al. (33) for an elegant treatment of this problem]. By contrast, because blur is a function of the microscope system and principally the objective lens, it can be modeled with relative simplicity. Such a model makes it possible to reverse the blurring process, and deconvolution uses this model to reverse or remove blur.

The Point-Spread Function

The model of blur that has evolved in theoretical optics is based on the concept of a point-spread function (PSF). This concept is very important to deconvolution and should be clearly understood to avoid imaging artifacts. Inoué and Spring (31) and Keller (36) provide good introductions to the concept of the PSF that are recommended for further detail. Several Web sites also provide tutorial information on the PSF, deconvolution, and 3-D microscopy generally. Two sites we recommend are http://www.microscopy.fsu.edu/primer/ and http://3Dmicroscopy.wustl.edu/~josec/tutorials/.

To understand the PSF, consider an infinitely small "point source" of light. Because the imaging system collects only a fraction of the light emitted by this point, it cannot focus this light into a perfect 3-D image of the point. Instead, the point appears widened and spread into a 3-D diffraction pattern. The 3-D diffraction pattern of an ideal point source of light is the PSF.

Depending on the imaging modality being used (wide-field, confocal, transmitted light), the PSF has a different shape. In a wide-field fluorescence microscope, the shape of the PSF is an oblong football of light surrounded by a flare of widening rings. To describe it in three dimensions, we apply a coordinate system of three axes (x, y, and z) where x and y are parallel to the focal plane of the specimen and z is parallel to the optical axis of the microscope. In this case, the PSF looks like a set of concentric rings in x-y, and it looks like an hourglass in x-z or y-z (Figure 1A). An x-y image through the center of the wide-field PSF looks like a set of concentric rings: this is the "Airy disk" of classical light microscopy.

How does the PSF affect image formation in a microscope? In the theoretical model of image formation, the PSF is considered the basic unit of any image. The PSF is to the image what the brick is to the house. The best an image can ever be is an assembly of PSFs, and increasing magnification will not change this. A noted theoretical optics textbook explains, "It is impossible to bring out detail not present in the primary image by increasing the power of the eyepiece, for each element of the primary image is a small diffraction pattern, and the actual image, as seen by the eyepiece, is only the ensemble of the magnified images of these patterns" (7).

As an example, consider a population of tiny fluorescent beads mounted under a cover slip. An in-focus image of this specimen will show a cloud of dots, each of which, when examined at high resolution, is actually a disk surrounded by a tiny set of rings (i.e., an Airy disk) (Figure 1B, 0 µm). If this specimen is brought out of focus slightly, a large set of concentric rings will appear where each dot was in the focused image (Figure 1B, 1 µm). If you collect a 3-D image of this specimen, then you will record a PSF at each bead. The PSF describes what happens to each point source of light after it passes through the imaging system.

The blurring process described above is mathematically modeled as a convolution. The convolution operation describes the application of the PSF to every point in the object: light emitted from each point in the object is convolved with the PSF to produce the image. This convolution causes points in the object to become blurred regions in the image. The brightness of every point in the image is linearly related by the convolution operation to the fluorescence of each point in the object (5,22). Since the PSF is 3-D, blurring from the PSF is an inherently 3-D phenomenon. The image from any focal plane contains blurred light from points located in that plane mixed with blurred light from points in other focal planes.

The situation can be summarized by saying that the image is formed by a convolution of the object with the PSF. Deconvolution reverses this process and attempts to reconstruct the object.

Aberrations in the PSF

The PSF can be defined either theoretically using a mathematical model of diffraction (20) or empirically by acquiring a 3-D image of a fluorescent bead (Figure 1A). A theoretical PSF generally has axial and radial symmetry [i.e., it is symmetric above and

below the x-y plane (axial symmetry) and rotationally about the z-axis (radial symmetry)]. An empirical PSF can deviate significantly from perfect symmetry (Figure 1A). This deviation, or aberration, is caused by irregularities or misalignments in any component of the imaging system light path, especially the objective lens but also other lenses, mirrors, filters, apertures, etc. The higher the quality of the optical components and the better the alignment, the closer the empirical PSF comes to its ideal symmetrical shape. Both confocal and deconvolution microscopy depend on the PSF being as close to the ideal case as possible.

Keller (36) provides a survey of all known types of PSF aberration. The most common type of aberration, well known to any professional microscopist, is spherical aberration. This is an axial asymmetry in the shape of the PSF, with a corresponding increase in size, particularly along the z-axis (Figure 1A). The result is a considerable loss of resolution and signal intensity (25). In practice, the most common cause of spherical aberration is a mismatch between the refractive indices of the lens immersion medium and the mounting medium in which the specimen rests (25,26). We emphasize the importance of minimizing this omnipresent aberration. While deconvolution can partially restore lost resolution (45), no amount of image processing can restore lost signal. See the "Artifacts and Aberrations" section for further discussion of this problem.

VARIETIES OF DECONVOLUTION ALGORITHMS

Now that we have surveyed a small amount of optical theory, we proceed to the business of this article, namely deconvolution algorithms—what they are and how they work. We will not review these algorithms in detail since published works already do so (32,37,48, 50,53), but we will explain their basic principles. We divide the available deconvolution algorithms into two classes, "deblurring" and "image restoration". Deblurring algorithms are fundamentally 2-D, because they apply an operation plane-by-plane to each 2-D plane of a 3-D image stack. In contrast, image restoration algorithms are properly 3-D because they operate simultaneously on every pixel in a 3-D image stack.

A few more technical terms must be defined. The *object* refers to the 3-D pattern of light emitted by fluorescent structures in the microscope's field of view. The *raw image* refers to an unprocessed digital image or image stack acquired from the microscope. Particular regions within the image are referred to as *features*.

Deblurring Algorithms

The algorithms called "nearest-neighbor", "multi-neighbor", "no-neigh-

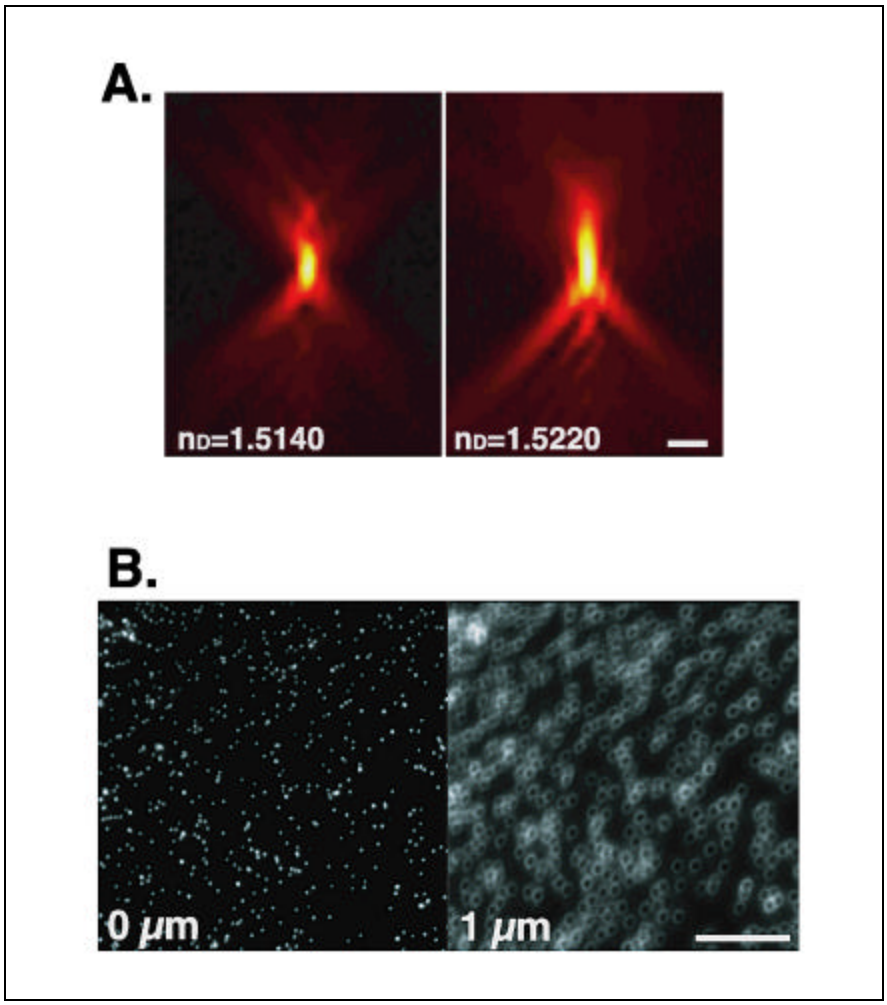

Figure 1. Effect of spherical aberration on the PSF. (A) x-z projections of two PSFs showing different degrees of spherical aberration. The optical axis is parallel to the vertical axis of the image. Left: shows minimal spherical aberration. Right: shows significant spherical aberration. (Note the axial asymmetry and widening of the central node along the optical axis in the right-hand image. This leads to degraded axial resolution and blurring of signal.) In theory, the size of the PSF is infinite, and the total summed intensity of light in planes far from focus is equal to the summed intensity at focus. However, light intensity falls off quickly and eventually becomes indistinguishable from noise. In an unaberrated PSF recorded with an NA 1.4 oil lens (a typical high-resolution lens), light occupying 0.2 μm^2 at the plane of focus is spread out over 90 times that area at 1 μm above or below focus. In these images, the object is a 0.1-μm (subresolution) fluorescent bead (Molecular Probes, Eugene, OR, USA) mounted in glycerol (RI = 1.47), with immersion oil of RI = 1.5140 (left) and 1.5220 (right). The images are sections from 3-D image stacks collected using a DeltaVision® wide-field imaging system (Applied Precision, Issaquah, WA, USA). Objective lens 100×/NA 1.4. Excitation wavelength, 580 nm; emission wavelength, 630 nm. Scale bar, 1 μm. (B) A population of 0.1-μm beads adhering to the surface of a cover slip, in focus (0 μm), and at 1 μm defocus (1 μm). Imaging parameters are as in panel A. The two images are scaled differently to emphasize the out-of-focus rings in the defocused image. Images are used by permission of ©Macmillan Press, 1998.

bor", or "unsharp masking" are fundamentally 2-D. We refer to them here as "deblurring algorithms". These algorithms apply an operation plane-by-plane to each 2-D plane of a 3-D image stack. For example, the nearest-neighbor algorithm operates on a plane z by blurring its neighboring planes $z \pm 1$ (using a digital blurring filter), then subtracting the blurred planes from z (1,10,48,50). Multi-neighbor methods extend this concept to a user-selectable number of planes. A 3-D stack is processed by applying the algorithm to every plane in the stack. In this way, an estimate of the blur is removed from each plane.

The deblurring algorithms are computationally economical because they involve relatively simple calculations on single image planes. However, there are major disadvantages to these approaches. First, noise from several planes is added together. Second, deblurring algorithms remove blurred signal and thus reduce signal levels. Third, features whose PSFs overlap in z may be sharpened in planes where they do not really belong (i.e., the apparent position of features may be altered). This problem is particularly severe when deblurring single 2-D images because they can contain diffraction rings or light from other structures that will then be sharpened as if they were in that focal plane. Taken together, these findings mean that deblurring algorithms improve contrast, but they do so at the expense of decreasing signal-to-noise ratio and may also introduce structural artifacts.

Two-dimensional deblurring algorithms may be useful in situations when a quick deblurring is needed or when computer power is limited. They work best on specimens that have fluorescent structures distributed discretely, especially in the z-axis. However, these algorithms cause artifactual changes in the relative intensities of pixels; therefore, morphometric measurements, quantitative fluorescence intensity measurements, and intensity ratio calculations should never be performed after applying a 2-D deblurring algorithm.

Image Restoration Algorithms

Image restoration algorithms deal with blur as a 3-D problem. Instead of "subtracting" blur, they attempt to "reassign" blurred light to an in-focus location. This is done by reversing the convolution operation inherent in the imaging system. If the imaging system is modeled as a convolution of the object with the PSF, then a deconvolution of the raw image should restore the object. However, the object cannot be restored perfectly because of the fundamental limitations inherent in the imaging system and the image-formation model (1). The best we can do is to estimate the object given these limitations. Restoration algorithms estimate the object, following the logic that a good estimate of the object is one that, when convolved with the PSF, gives back the raw image.

An advantage of this formulation is that convolution operations on large matrices (such as a 3-D image stack) can be computed very simply using the mathematical technique of Fourier transformation. If the image and PSF are transformed into "Fourier space", the convolution of the image by the PSF can be computed simply by multiplying their Fourier transforms. The resulting Fourier image can then be back-transformed into real 3-D coordinates. (For an introduction to Fourier transforms in optics, see Reference 22.)

Inverse Filters

The first image deconvolution algorithms to be developed were "inverse filters". Such filters, along with their cousins the "regularized inverse filters", have been used in electronic signal processing since the 1960s and were applied to images in the 1970s (10). In image-processing software, these algorithms go by a variety of names including "Wiener deconvolution", "Regularized Least Squares", "Linear Least Squares," and "Tikhonov-Miller regularization" (10,37,48,55).

An inverse filter takes the Fourier transform of an image and divides it by the Fourier transform of the PSF. Since division in Fourier space is equivalent to deconvolution in real space, this is the simplest way to reverse the convolution that produced the blurry image. The calculation is rapid, about as fast as the 2-D deblurring methods discussed above. However, the utility of this method is limited by noise amplification. During division in Fourier space, small noise variations in the Fourier transform are amplified by the division operation. The result is that blur removal is traded against a gain in noise. Also, an artifact known as "ringing" can be introduced (see "Aberrations and Artifacts" section).

Noise amplification and ringing can be reduced by making some assumptions about the structure of the object that gave rise to the image. For instance, if we assume that the object was relatively smooth, we can eliminate noisy solutions with rough edges. This approach is called regularization. A regularized inverse filter can be thought of as a statistical estimator that applies a certain kind of "constraint" on possible estimates, given some assumption about the object: in this case, smoothness. A constraint on smoothness enables the algorithm to select a reasonable estimate out of the large number of possible estimates that can arise because of noise variability.

Regularization can be applied in one step within an inverse filter (42), or it can be applied iteratively. The result is usually smoothed (i.e., stripped of higher Fourier frequencies). Much of the "roughness" being removed here occurs at Fourier frequencies well beyond the resolution limit and, therefore, does not eliminate structures recorded by the microscope. However, since there is a potential for loss of detail, software implementations of inverse filters typically include an adjustable parameter that allows the user to control the tradeoff between smoothing and noise amplification (37).

Constrained Iterative Algorithms

To improve the performance of inverse filters, a number of other 3-D algorithms can be used to restore images (48,50,53). These methods are called "constrained iterative algorithms". They work in successive cycles and are therefore called "iterative". They also usually apply constraints on possible solutions. These constraints not only help to minimize noise or other distortion but also increase the power to restore blurred signal.

A typical constrained iterative algorithm works as follows. An estimate of

the object is made (this is usually the raw image itself). The estimate is convolved with the PSF, and the resulting "blurred estimate" is compared with the raw image. This comparison is used to compute an error criterion that represents how similar the blurred estimate is to the raw image. This error criterion (or "figure of merit") is then used to alter the estimate in such a way that the error is reduced. A new iteration then takes place: the new estimate is convolved with the PSF, a new error criterion is computed, etc. The best estimate will be the one that minimizes the error criterion: therefore, as long as the error criterion has not been minimized, each new estimate is blurred again, an error criterion is computed, etc. This process is repeated until the error criterion is minimized or reaches a defined threshold. The final restored image is the object estimate at the last iteration.

Most algorithms incorporate constraints on the range of allowable estimates. One commonly used constraint is smoothing or regularization, as discussed above. As iterations proceed, the algorithm will tend to amplify noise, so most implementations suppress this with a smoothing or regularization filter. Another common constraint is "nonnegativity" (1). This means that any pixel value in the estimate that becomes negative during the course of an iteration is automatically set to zero. (Pixel values can become negative, either because of Fourier transformation or a subtraction operation in the algorithm.) The nonnegativity constraint is realistic because an object cannot have negative fluorescence. It is essentially a constraint on possible estimates given our knowledge of the object's structure. Other types of constraints include "boundary constraints" on pixel saturation, constraints on noise statistics, and other statistical constraints.

Classical algorithms. The first applications of constrained iterative deconvolution algorithms to microscope images were based on the Jansson-Van Cittert (JVC) algorithm, a procedure developed by Van Cittert (19) for use in spectroscopy and adapted by Jansson (32). This algorithm was modified by Agard for application to digital microscope images (1,21). Various implementations of Agard's modified algorithm are currently marketed by Vaytek, Intelligent Imaging Innovations, Applied Precision, Carl Zeiss, and Bitplane. In addition, Carrington and co-workers developed a regularized least squares minimization method (8,17) that has been marketed by Vaytek and Scanalytics. These algorithms use an additive or multiplicative error criterion to update the estimate at each iteration (37,48,50,53).

Statistical algorithms. Another family of iterative algorithms uses probabilistic error criteria taken from statistical theory (12,27,48). "Likelihood", a kind of reverse of probability (16), is used in the maximum likelihood estimation (MLE) and expectation maximization (EM) algorithms implemented by SVI, Bitplane, ImproVision, Carl Zeiss, and Autoquant. MLE is a popular statistical tool with applications in many branches of science. A related statistical measure, maximum entropy (ME—not to be confused with EM) has been implemented in image deconvolution by Carl Zeiss.

Statistical algorithms are more computationally intensive than the classical methods and can take significantly longer to reach a solution (see Reference 48 for an estimate). However, they may restore to a slightly higher degree of resolution than the classical algorithms. They also have the advantage that they impose constraints on the expected noise statistic (i.e., a Poisson or a Gaussian distribution). As a result, these algorithms have a more subtle noise policy than simply regularization, and they may give better results on noisy images. However, the choice of an appropriate noise statistic may depend on the imaging condition, and some commercial software packages are more flexible than others in this regard.

Blind deconvolution. Blind deconvolution is a relatively new method that greatly simplifies the use of deconvolution for the non-specialist. Currently, only Autoquant implements this method. It was developed by altering the MLE algorithm so that not only the object but also the PSF is estimated (27). In this approach, an estimate of the object is made. This estimate is convolved with a theoretical PSF calculated from optical parameters of the imaging system. The resulting blurred estimate is compared with the raw image, a correction is computed, and this correction is used to generate a new estimate (as explained above). This same correction is also applied to the PSF, generating a PSF estimate. In further iterations, the PSF estimate and the object estimate are updated together.

Blind deconvolution works well not only on high-quality images but also on noisy or spherically aberrated images. It begins with a theoretical PSF but adapts it to the specific data being deconvolved. In this regard, it spares the user from the difficult process of acquiring a high-quality empirical PSF (26,37). Also, because it adjusts the PSF to the data, it can partially correct for spherical aberration. However, this computational correction should be a last resort—it is preferable to minimize spherical aberration during image acquisition (see "Aberrations and Artifacts" section).

Deconvolution of confocal and multi-photon images. As one might expect, it is also possible to restore confocal or multi-photon microscope images. The combination of confocal microscopy and deconvolution improves resolution beyond what is attainable with either technique alone (49). However, the major benefit of deconvolving a confocal image is not so much the reassignment as the averaging of out-of-focus light. This results in decreased noise (e.g., see Reference 13). Deconvolution of multiphoton images has also been used to successfully remove image artifacts and improve contrast (52). In all of these cases, care must be taken to use the appropriate PSF, especially if the confocal pinhole is adjustable. The interested reader is referred to a previous discussion on the implementation of deconvolution algorithms for laser-scanning microscopes (28).

Implementation. Processing speed and quality is dramatically affected by how a given deconvolution algorithm is implemented in software. The algorithm can be implemented in ways that reduce the number of iterations and accelerate convergence to a stable estimate. For example, the unoptimized JVC algorithm usually requires 50–100 iterations to converge to an optimal estimate (2,48,50). By prefiltering the raw image to suppress noise and cor-

recting with an additional error criterion on the first two iterations, the algorithm converges in only 5–10 iterations. In addition, a smoothing filter is usually introduced every five iterations to curtail noise amplification.

When using an empirical PSF, it is critical to use a high-quality PSF with minimal noise. No deconvolution package we know of uses the "raw" PSF recorded from the microscope. Instead, the packages contain preprocessing routines that reduce noise and enforce radial symmetry by averaging the Fourier transform of the PSF. Many packages also enforce axial symmetry in the PSF and thus assume the absence of spherical aberration. These steps reduce noise and aberrations and make a large difference in the quality of restoration.

Another aspect of implementation is preprocessing of the raw image, via routines such as background subtraction, flatfield correction, bleaching correction, lamp jitter correction, etc. These operations can improve the signal-to-noise ratio and remove certain kinds of artifacts. Most available packages include such operations, and the users' manuals of the packages should explain them. We will return to them below in our discussion of artifacts.

Other implementation issues concern data representation. Images can be divided into subvolumes or represented as whole data chunks. Individual pixel values can be represented as integers or as floating-point numbers. Fourier transforms can be represented as floating-point numbers or as complex numbers. In general, the more faithful the data representation, the more memory and processor time it requires. Thus, there is a tradeoff between the speed of computation and the quality of restoration. These issues are discussed below in the "Artifacts and Aberrations" section.

Summary. Iterative restoration algorithms differ from both deblurring algorithms and confocal microscopy in that they do not remove out-of-focus blur but instead reassign it. In this way, out-of-focus signal is used rather than thrown away. After restoration, pixel intensities within fluorescent structures increase. However, the total summed intensity of each image stack stays the same, as intensities in formerly blurred areas diminish. Blur surrounding details of the object is moved back into focus, resulting in sharper definition of object from background, better contrast, and improved signal-to-noise ratio.

These properties are shown in Figure 2. Restoration improves image contrast and subsequently allows better resolution of objects, without the introduction of noise that occurs in deblurring methods (Figure 2A). Perhaps more importantly for image analysis and quantitation, the sum of the fluorescence signal in the raw image is identical to that in the deconvolved image. When properly implemented, image restoration methods preserve total signal intensity but improve contrast by adjustment of signal position (Figure 2B). Therefore, quantitative analysis of restored images is possible and, because of the improved contrast, often desirable.

When used in conjunction with wide-field microscopy, iterative restoration methods are light efficient. This is most valuable in light-limited applications such as high-resolution fluorescence imaging, where objects are typically small and contain few fluorophores (15,18), or in live-cell fluorescence imaging, where exposure times are limited by the extreme sensitivity of live cells to phototoxicity (9,24,46,54).

PERFORMANCE ISSUES

Now that we have explained the principles on which deconvolution algorithms are based, we can offer some technical insights on how they perform and how best to compare them.

Resolution and Contrast Improvement

What kind of quantitative improvement in image quality can be expected from iterative deconvolution? We have attempted to answer this by measuring the size and brightness, before and after deconvolution, of a test object of known size (Figure 3). The object is a 0.1-μm (subresolution) fluorescent bead. This is a nearly ideal specimen—there is no out-of-focus signal coming from any other object, and all aberrations were carefully minimized before data collection. In the raw image, the bead measures 0.7 μm along the z-axis; in the deconvolved image, it measures 0.45 μm (Figure 3A). This is a modest improvement in resolution, which may reveal a biologically interesting structure in only a limited number of cases.

However, the major change in the image is shown in Figure 3B. This plot shows the integrated pixel intensity—the sum of all pixel values in each focal plane—as a function of focal depth. There is significant out-of-focus intensity at 2 μm away from the bead before processing. Restoration by iterative deconvolution moves the majority of the out-of-focus intensity back to its focal plane of origin. The result is a significant improvement in image contrast, making it easier to resolve and distinguish features in the image.

Comparisons between Algorithms

Which iterative deconvolution algorithm gives the best restoration? A number of Web sites compare the results of different algorithms, but these comparisons can be misleading for a variety of reasons. First, algorithms are often compared using images of synthetic spherical objects (such as the bead in Figure 3) or even computer-generated images of theoretical objects. The relationship between performance with such test objects and performance with real biological specimens is not straightforward. Furthermore, unless the comparison is done quantitatively, with objects of known size, it is hard to know whether a more pleasing result is really more accurate. For instance, the algorithm might eat away the edges of features, making them look sharper but confounding measurements.

In addition, algorithm comparisons are usually published by biased parties with an interest in the result of the comparison. Frequently, these parties compare an algorithm whose implementation they have developed and optimized over many years, with a non-optimized algorithm they have implemented straight out of the book. However, as noted above, big differences in speed, stability, and resolution improvement can be attributed to implementation and optimization of the algorithm. Therefore, the only fair comparison is between realized software packages, rather than between algorithms.

We recommend that you, the potential buyer of deconvolution software, should compare the performance of various software packages on your own data. Unfortunately, this may require a certain amount of determination. To compare images restored by different packages, you will need to make sure the sales representative saves your images in a consensus file format because, otherwise, each software company will save to its own proprietary file format, which often cannot be opened by competitors' software. Currently, the format with the widest circulation is a stack of sequentially numbered TIFF files, each representing a focal plane (not a projection!) of the 3-D image. You will also need a few example image files and PSF files that you can run through different packages. If you do not have experience collecting PSF files, you may want to gain some before testing deconvolution software that requires an empirical PSF. It is better to test with your own images and PSFs because the preparation you use, along with the lens, magnification, noise level, signal intensity, spherical aberration, etc., will affect the quality of deconvolution tremendously.

Speed and Memory Usage

Increasing processor speed, RAM, and bus speed all increase the speed of deconvolution. During deconvolution, a number of large arrays representing different forms of the image are stored simultaneously in RAM and are moved around inside the computer via the bus. As a result, RAM is critical for the rapid processing of 3-D images, as is bus speed. As a rule of thumb, your computer should have at least three times as much RAM as the size of the image you wish to deconvolve. Also, computers with fast buses perform much better, even with nominally slower processors.

The size of an image file is usually reported by the operating system. However, if in doubt, it can be calculated by multiplying the total number of pixels in the image by the number of bytes/pixel ("bit depth"). The bit depth is originally set by the camera, which may produce 8, 10, 12, or 16 bits/pixel. Once the image is acquired, the bit depth is determined by your software and computer system (almost always it will be 8 or 16 bit; 8 bits = 1 byte). In a multicolor image, each color must be stored and deconvolved separately, so one must be careful to get the bit depth for the whole image, not just for one color channel. An example: a 3-D stack, where each plane is 512 x 512 pixels, containing 64 optical planes, with three colors at 8 bits/pixel (= 1 byte/pixel) measures $512 \times 512 \times 64 \times 3 \times 1 = 50$ MB. The image file header may add slightly to this size.

ARTIFACTS AND ABERRATIONS

After deconvolution, the restored image may include apparent artifacts (e.g., striping, ringing, or discontinuous cytoskeletal staining). Sometimes these problems are related to data representation and will not occur with a different algorithm or software package. They can also occur when processing parameters are not set appropriately for the raw image. Finally, artifacts are often not caused by computation, but by his-

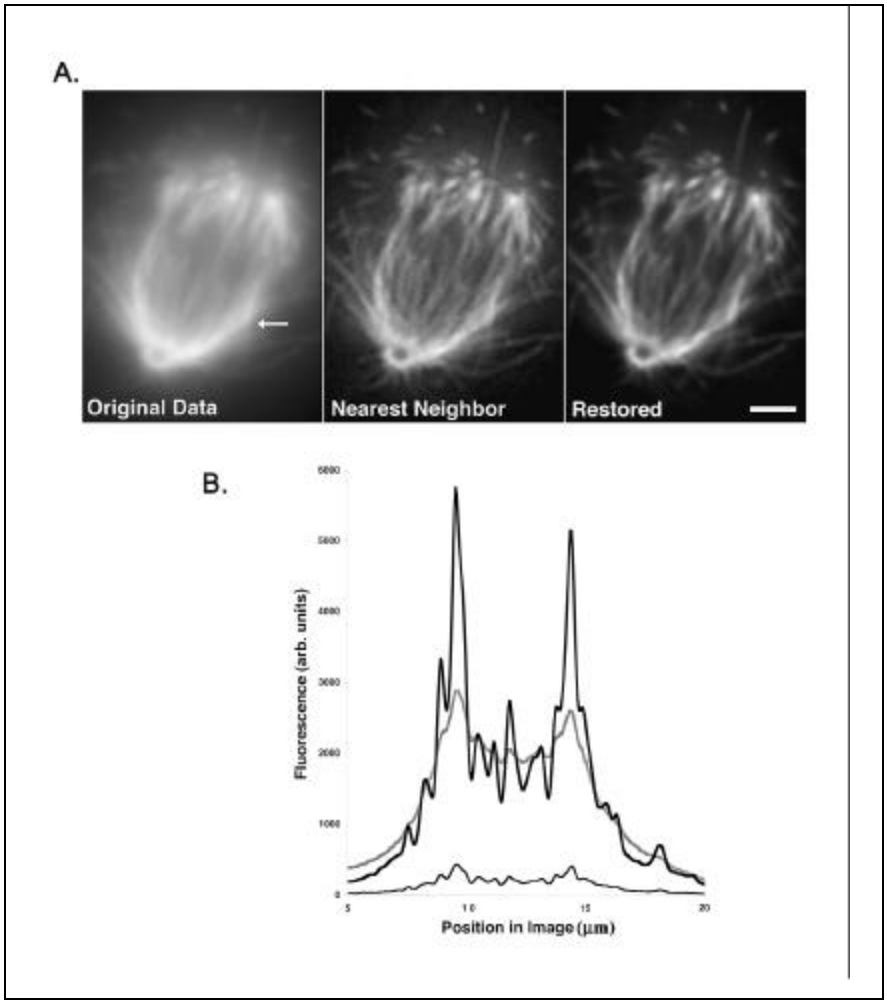

Figure 2. Comparison of deblurring and restoration methods. Data from a 3-D image stack containing 70 optical sections, each separated by 0.2 μm, recorded using same apparatus as in Figure 1. The object is a XLK2 cell fixed with 3.7% formaldehyde, stained with mouse anti-tubulin and Cy5-conjugated donkey anti-mouse IgG.) (A) A single focal plane from the 3-D stack is shown before any processing (Original Data), after deblurring by a nearest-neighbor algorithm (Nearest Neighbor), and after restoration by constrained iterative deconvolution (Restored) using the DeltaVision softWoRx™) (Applied Precision) deconvolution software. Both deblurring and restoration improve contrast, but the signal-to-noise ratio is significantly lower in the deblurred image than in the restored image. Scale bar, 2 μm. Arrow shows the position of the line plot presented in panel B. (B) Plot of pixel brightness values along a horizontal line shown by the arrow in panel A. Original data (gray line), deblurred (thin black line), restored (thick black line). Deblurring (or any other 2-D filter) causes a significant loss of pixel intensity all across the image, whereas restoration causes a gain of intensity in areas of detail.

tology, optics, or electronics. When trying to diagnose the cause of an artifact, the first step to take is a careful comparison of the raw image with the deconvolved image. If the artifact is visible in the raw image, then it must be caused by factors upstream from deconvolution (i.e., by specimen preparation, optics, or electronics). By adjusting the contrast and brightness of the raw image, you can sometimes detect an artifact you would not have noticed initially. If the artifact is not detectable in the raw image, then some aspect of deconvolution is implicated. In this case, it may be useful to compare the results of deconvolution by different kinds of algorithms [e.g., an inverse filter versus a constrained iterative algorithm (35)].

The PSF

The quality of the PSF is critical to the performance of a deconvolution algorithm. A noisy, aberrated, or improperly scaled PSF will have a disproportionate effect on the results of deconvolution. This is especially true for the iterative methods because the PSF is repeatedly applied. In all cases, the distribution and extent of blurred light in the raw image must match the PSF. If a mismatched PSF is used, then artifacts may result or restoration quality may diminish. In the next sections, we discuss common problems with the PSF and ways to correct them.

Theoretical versus Empirical PSFs

In many deconvolution packages, the user can choose either a theoretical or an empirical PSF for image restoration. In general, results are better if an empirical PSF is used. Procedures for acquiring an empirical PSF can be found in (26,37). There are several reasons not to preferentially use a theoretical PSF. First, although good theoretical models for the PSF exist (20), they are not perfect models, and an empirical PSF contains information not available in theoretical models. Second, the theoretical PSF available in commercial software packages generally assumes perfect axial and rotational symmetry: this means it may misfit the distribution of blur in the raw image. This problem is most serious at high resolution [e.g.,

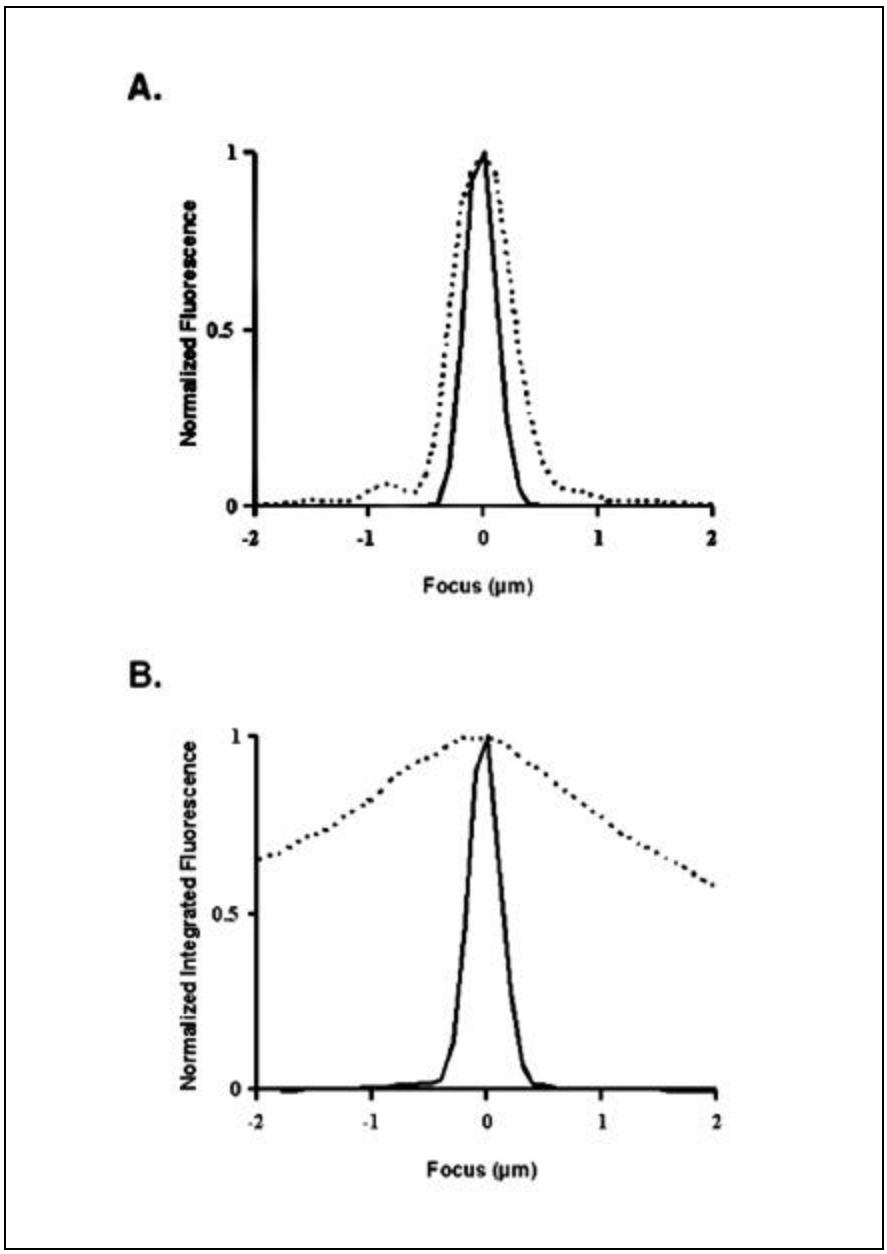

Figure 3. Restoration significantly improves contrast and modestly improves resolution. Data from image stacks of a subresolution fluorescent bead (sample and data collection as in Figure 1A, left). Dotted line, raw image. Solid line, image restored using constrained iterative deconvolution as in Figure 2. For graphical purposes, each intensity value has been normalized to the maximum value of its own image stack. Without such normalization, data from the raw image would barely be visible on the graph because pixel intensities near focus are so much brighter in the restored image than in the raw image. (Nonetheless, the total integrated intensity of each image stack is the same in both cases.) (A) To quantify the improvement in resolution, we measured the pixel intensities on a line parallel to the optical axis through the middle of the bead, before and after restoration. Each pixel intensity is normalized to the minimum and maximum values of its own image stack. Pixel intensity is plotted as a function of distance along the z-axis from the center of the bead (0 µm). The FWHM of the z-axis intensity profile is 0.7 µm in the raw image and 0.45 µm in the restored image. (The actual object measures 0.1 µm.) This modest increase in resolution in the restored image will only rarely reveal a biological structure that was not visible in the raw image. (B) To quantify the improvement in contrast, we summed all pixel intensities in each focal plane of the raw and restored image stacks. The summed intensity of each plane is normalized to the minimum and maximum values of its image stack for comparison. Summed intensity is plotted as a function of focal (z-axis) distance from the center of the bead (0 µm). Restoration causes a movement of signal intensity from out-of-focus volume to in-focus, resulting in a major improvement in contrast and signal-to-noise ratio. However, the integrated intensity of the whole image (i.e., the sum of the summed intensities of each plane) is the same in the raw and restored images.

numerical aperture (NA) 1.2–1.4], where small manufacturing variations in the lens can cause minor aberrations in the symmetry of the PSF (36,45). A theoretical PSF does not reflect these lens-specific variations and yields inferior deconvolution results. Furthermore, an empirical PSF can be helpful in selecting a lens. Aberrations that are hard to detect when looking at complicated objects are very clear when examining the PSF from a single bead. Therefore, before purchasing a costly new objective lens, we recommend that you acquire PSFs from several lenses and choose the one with the most ideal PSF. Third, an empirical PSF allows you to measure the performance of your imaging system. Many potential problems that may occur during an experiment (e.g., stage drift, lamp flicker, camera noise, refractive index (RI) mismatches, temperature changes due to heavy ventilation, etc.) will also occur during PSF acquisition and be more easily discernible. Therefore, aberrations in the empirical PSF suggest ways to improve your microscopy.

When acquiring an empirical PSF, care must be taken to match the aberrations of the raw image. Ideally, both the raw image and the PSF should be free of aberration, but this is not always possible. If major aberrations are present in the raw image, then they should if possible be matched by aberrations in the PSF (45). Otherwise, the deconvolved image may contain errors or be poorly restored. In addition, if the PSF is noisy, then substantial noise will appear in the deconvolved image. To reduce noise and eliminate minor aberrations, many packages radially average the PSF or average the images of several beads to create a smoother PSF. In addition, most commercial packages automatically interpolate the PSF sampling interval to match the sampling interval of the raw image. If this is not the case, the PSF and raw image must be acquired at the same sampling interval.

When using a theoretical PSF, the PSF parameters must be set appropriately. The PSF parameters are imaging modality, NA, emitted light wavelength or λ, pixel size, and z-step. These parameters affect the size and shape of the PSF. In general, PSF size increases with increasing wavelength and with decreasing NA. (See Appendix on resolution criteria.) The pixel size and z-step parameters are used for scaling the PSF with respect to the raw image.

If the size and shape of the theoretical PSF are not appropriate to the raw image, then artifacts can result for several reasons. First, the algorithm interprets the sampling interval of the image in terms of the PSF size. Second, the PSF determines the size and shape of the volume from which blurred light is reassigned. If this volume does not correspond to the true distribution of blur in the image, then artifacts result. This can happen if there is mis-scaling or aberration of the PSF. It can also occur if the "real PSF" of the image has an aberration that is not matched by the algorithm's PSF.

When using a theoretical PSF, there are cases where, paradoxically, a better result may be obtained if the PSF is too large. A possible explanation is that the real PSF of the raw image is larger than expected. This might occur because of refractive index mismatch, which causes spherical aberration and z-axis scaling. Both phenomena widen the real PSF in z, making it larger and thereby lowering the effective NA of the lens. If this is suspected, then try subtracting a small increment (e.g., 0.05) from the NA of the theoretical PSF. In some software, a similar result is obtained by setting the z-step of the PSF to be smaller than the z-step of the raw image.

Normally, the z-step of the PSF should always be identical to the z-step of the raw image: this ensures that the scaling of the PSF is appropriate to the imaging conditions. With an empirical PSF, however, it may be possible to acquire a PSF at finer z-resolution than the raw image. This makes the Fourier transform of the PSF more detailed and can benefit restoration. However, this trick works only if the software can interpolate the image's sampling interval and make it correspond to the PSF. Consult your user manual to determine how your software handles the PSF.

Spatial Variation of the PSF

Currently, most commercially available deconvolution packages assume that the PSF is constant for all points in the object, a property known as spatial invariance. Microscope optics generally meet this assumption; however, other issues such as refractive index gradients in the specimen material or mismatch of immersion and mounting media, cause spatial variations in the PSF, especially in thick specimens. At present, all commercial software packages assume spatial invariance. However, increasing computer power may make it feasible to vary the PSF through the image in the near future. It may also become possible to correct spatial variations in the PSF by using a transmitted light image to map refractive index gradients in the specimen and adjust the PSF accordingly (33).

Spherical Aberration

A notorious kind of PSF aberration, and one of the most difficult to combat, is spherical aberration. It involves an axial asymmetry in the shape of the PSF, which both increases the "flare" of the PSF and decreases its brightness. This is one of the primary causes of degraded resolution and signal loss in both confocal and wide-field microscopy (25,26). Spherical aberration can be detected by focusing up and down through the specimen, looking for asymmetry in the out-of-focus rings above and below a brightly fluorescent point-like detail (Figure 4, A–D). Alternatively, it can be detected in an acquired image stack: when viewing the stack in x-z or y-z projection, look for axial asymmetry in the flare of blurred light around a fluorescent structure. A third way to detect spherical aberration is to acquire a PSF image from a fluorescent bead mounted under similar optical conditions as the specimen and to look for axial asymmetry in the flare of the PSF.

If the specimen is comparatively thick (>10 μm), spherical aberration may be induced gradually as you image deeper into the specimen. Therefore, a bead mounted directly under the cover slip surface may not reveal spherical aberration. For this reason, some people recommend acquiring the PSF image from a bead located within a piece of tissue (e.g., by soaking the tissue in a solution of fluorescent beads). However, this PSF should not be used for deconvolution because scatter from the

tissue will make the PSF noisier.

Spherical aberration is caused by an imperfection in the light path of the imaging system. This can be due to defects in the objective lens, but more frequently by mismatches of refractive index (RI) in the optical media in front of the lens. Objective lenses are usually corrected to minimize spherical aberration but only if they are used with the proper type of cover slip glass, the proper cover slip thickness, and the proper immersion and mounting media. The optical properties of these materials are essential for the proper focusing of light by the objective lens.

Quite commonly in biology, the RI of the immersion and mounting media are not the same. In these cases, two distortions may occur: spherical aberration and a scaling of z-axis distances. Hell et al. (25) describe both phenomena in detail. Both phenomena depend on focal depth, so features at different depths will show different amounts of spherical aberration and z-scaling. Z-scaling does not affect resolution or signal intensity: it is simply a linear scaling of z-axis distance measurements by the ratio of the RIs of the mismatched media (25). To correct for this distortion, z-distances can be multiplied by a scalar compensation factor, and some software packages offer this feature. On the other hand, spherical aberration is difficult to correct and, therefore, may be tempting to ignore. However, a little attention to this issue can bring big improvements in image quality, especially under low-light conditions such as in living cells.

There has been interest in digitally correcting for spherical aberration by deconvolving with a spherically aberrated PSF (45). This requires deconvolution software that does not automatically preprocess the PSF to make it axially symmetric. If this is the case, then one can precisely match the aberration in the image and PSF by having on hand a "family" of empirical PSFs with different degrees of spherical aberration and selecting the most appropriate PSF for a given imaging condition. This type of computational correction may restore lost resolution to some extent, but it cannot restore lost signal. Therefore, a better way to correct for spherical aberration is to eliminate it beforehand by optical means.

The following are optical methods to correct for spherical aberration. (*i*) Using a "dipping" objective (i.e., one that goes directly into the mounting medium without a cover slip). This way the immersion medium and the mounting medium are one, and no RI mismatch is possible. Distortion can still occur, however, because of RI mismatch between the specimen and the immersion medium or RI gradients within the specimen itself. (*ii*) Adjusting the RI of the immersion medium to compensate for the RI of the mounting medium (26). If the specimen is mounted in a medium of *lower* RI than glass (e.g., any glycerol- or water-based media), then you should *increase* the RI of the immersion medium. This method allows a relatively aberration-free imaging condition up to a limited focal depth (10–15 µm from the cover slip for a NA 1.4 lens observing a specimen mounted in glycerol). The RIs of a few common mounting media are listed by Bacallao et al. (4). Cargille Laboratories (Cedar Grove, NJ, USA) supplies a line of "Laser Liquid" immersion media with specified RIs. These can be mixed to any desired intermediate value. (*iii*) Using a specially corrected objective lens, such as the high-NA water-immersion lenses now available for use with cover-slipped specimens. These lenses have a correction collar that compensates for RI variations between the lens and the specimen. These lenses are expensive (currently about $10 000). However, they may provide the only way to eliminate spherical aberration when imaging deeper than 15 µm into a thick specimen. The correction collar on such lenses should not be confused with the variable-aperture correction collars found on less expensive dark-field lenses.

The Sampling Interval

The proper sampling interval in x, y, and z is important for good deconvolution results. The standard practice is to sample twice per resolvable element: this conforms to the Nyquist sampling theorem, which states that two samples per resolvable element are required for accurate detection of a signal (22). However, the Nyquist sampling frequency is really just the minimum necessary for a reasonable approximation of the real signal by discrete sampling. A higher sampling frequency gives better restoration, especially when using 3-D algorithms. In contrast, 2-D deblurring algorithms work best when the sampling rate is lower in z (i.e., when the spacing between optical sections is equal to or greater than the resolvable element).

For fluorescence microscopy the resolvable element is often defined using the Rayleigh criterion (see Appendix). For example, with the dye FITC (emitting at 520 nm), an NA 1.4 oil lens, and a mounting/immersion medium with an RI of 1.51, the resolvable element is 227 nm in x-y and 801 nm in z (according to the Rayleigh criterion). To sample at the Nyquist frequency, the sampling interval should be twice this, or 0.114 µm in x-y and 0.4 µm in z. To oversample sufficiently, we routinely sample at 0.07 µm in x-y and 0.2 µm in z under these conditions.

Of course, these guidelines represent optimal settings and should be balanced by considering the specific preparation in question. For example, the fluorescence signal may be so low (e.g., in a live cell experiment) that binning of CCD pixels is required, or events may occur so quickly that there is not time to finely sample in z. In such cases, a suboptimal sampling interval will be required. Fortunately, restoration algorithms still work quite well under these conditions (9,54), although some changes in regularization may be necessary in some algorithms (47).

Ringing and Edge Artifacts

Ringing is an artifact chiefly found in deblurring or inversion methods, but it does sometimes occur with iterative methods. It has the appearance of dark and light ripples around bright features of an image (Figure 4, E and F). It can occur in z and in x-y: in z, it looks like a shadow in a deeper z-section, outlining a fluorescent structure.

Ringing is generally caused by the conversion of a discontinuous signal into or out of Fourier space (38). A number of related issues can produce signal discontinuities and therefore cause ringing. Discontinuities can occur at the edges of the image or of subvolumes of the image (Figure 4D) or even

at the edges of bright features (Figure 4, E and F). Discontinuities can also arise if the spatial sampling of the raw image or PSF is too coarse, if the image or PSF are noisy, or if the PSF size or shape is inappropriate to the image.

The surest way to avoid ringing is by proper windowing of the Fourier transform. This is an implementation issue that some software packages have not incorporated but is increasingly the norm. Ringing can also be avoided by using a finer sampling interval in the image or PSF, by smoothing the image or PSF, by careful matching of aberrations in the PSF with those of the image, or by adjustment of PSF parameters.

Ringing at edges of the image can be suppressed by simply cutting out the edges. Many algorithms cut out a "guard band" of 8–10 pixels or planes around the entire image. This operation requires blank space to be left above and below the fluorescent structures in the image. Blank space is advisable anyway, because blurred light from above and below an object can be reassigned and contribute to the signal for that object. However, the blank space can be artificially created, without loss of fidelity, by adding interpolated planes to the top and bottom of the image stack. The resulting cost in memory and processing time can be minimized by performing the interpolation in Fourier space (2).

Ringing at the edge of subvolumes is a similar issue (Figure 4D). Usually, the amount of subvolume overlap can be increased so as to remove the artifact. Alternatively, if enough RAM is installed, then the whole image can be processed as a single volume.

Noise Amplification or Excessive Smoothing

Noise amplification has been alluded to already as an artifact caused by deblurring and inversion algorithms. It also occurs with iterative algorithms as repeated convolution operations introduce high-frequency noise. The artifact appears as a distinctive mottling of the image that often is constant in every plane and particularly noticeable in background areas. It is usually suppressed by a smoothing filter (or regularization or "roughness" filter) (53). If excessive noise is observed in a decon-

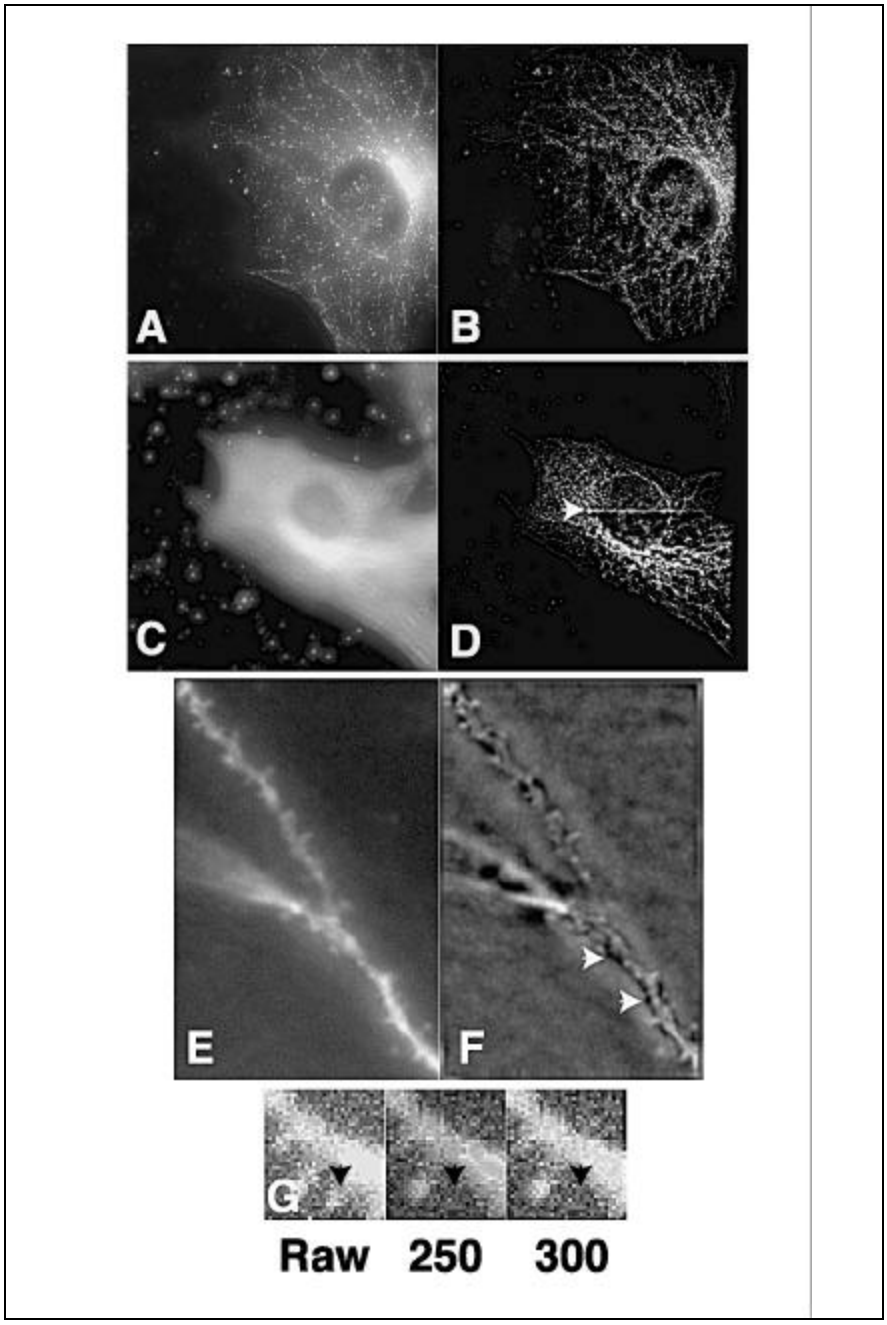

Figure 4. Examples of artifacts. (A–D) Images of bovine pulmonary endothelial cells from a prepared slide (Fluocells™; Molecular Probes). These cells were fixed, permeabilized, and stained with BODIPY FL™ phallicidin (emission max = 512 nm). Images were acquired as described in previous figures, varying the RI of the immersion oil: (A and B) RI = 1.514; (C and D) RI = 1.524. (A) Raw image (maximum-intensity projection) without appreciable spherical aberration. (B) The same image after deconvolution. (C) Raw image of a similar cell with significant spherical aberration (note diffraction rings around points surrounding the cell). (D) The same image as in panel C after restoration by iterative constrained deconvolution. Note that a horizontal line is introduced at the edge of a subvolume (arrowhead). This artifact is not seen in the raw image (C) and is, therefore, not caused by the imaging system. (E and F) Pyramid cell dendrites from sections of rat visual cortex, fixed and then injected with Lucifer Yellow (emission max = 533 nm). Each image is a single representative optical section from a 3-D stack. (E) Raw image. (F) After iterative deconvolution. In this case, significant ringing occurs (dark outline around fluorescent dendrite; arrowheads), probably because of spherical aberration in the raw image, which is not matched in the PSF. Note also noise amplification seen as a mottling of background in the deconvolved image. (G) Detail of a similar sample to panel E before deconvolution (Raw) and after deconvolution by increasing numbers of iterations (250 and 300). A small feature disappears (black arrow). This type of artifact can be avoided by correct use of PSF, as described in the text.

volved image, then the first step is to make sure the PSF is as noise-free as possible. If this is the case, the next step is to look at the filters in the algorithm. If the filters are adjustable by the user, then try to set the parameters such that image features larger than the resolution limit are maintained, while smaller variations are suppressed. Do not increase smoothing too much because excessive smoothing will degrade resolution and contrast.

Disappearing and Exploding Features

Perhaps the most annoying artifacts of deconvolution are an apparent loss of dim features or the "blowing up" of very bright ones. These artifacts differ from excessive smoothing in that they affect specific features, as opposed to the image as a whole (Figure 4G). In such a case, try to compare the effects of different algorithms on the appearance of the artifact (e.g., inverse filter vs. iterative algorithms, classical vs. statistical). If you see the artifact only with one algorithm, then the cause may be algorithm instability or a specific effect of the algorithm on a particular specimen type. If you see the artifact with several different algorithms, then there are a variety of possible causes. If small, dim features of the image are disappearing, then the cause can be a combination of nonnegativity and smoothing filters. This is likely if there is excessive noise in the image or the PSF or if conditions are favorable to ringing. Noise and ringing can cause negative pixel values so that the nonnegativity constraint is invoked and features are broken up. If this happens, then the smoothing filter may smooth away the remaining fragments, causing the feature to disappear. This can be avoided by reducing noise in the image and the PSF (e.g., with longer exposure times or by averaging), by eliminating PSF mismatches (as described above), and by properly adjusting the smoothing filter.

If bright features are expanding, then the cause is probably pixel saturation. When the raw image contains bright features, restoration will tend to increase their brightness even further. However, in very bright features, this can cause pixel values to exceed the

highest value representable at a given bit depth (i.e., these pixels will "saturate"—they will be assigned the maximum value). The feature will then become very bright, will appear to flatten out, and may expand in size. This can be avoided by using a camera and software that allow greater bit depth and an algorithm that internally stores data as floating point numbers. This increases memory demands but prevents the occurrence of such artifacts.

Specimen-Dependent Artifacts

There have been reports, both informal and published, of deconvolution artifacts seen with some kinds of specimens but not with others. In particular, Karpova et al. have reported that filamentous structures are broken up with one class of deconvolution algorithms but not with another (35,37). If a structural change is observed, then it is always advisable to compare the deconvolved and raw images. If close inspection reveals that a structure evident after deconvolution does not occur in the raw image, then make sure that the PSF parameters are set correctly and that the PSF is not noisy. If the problem persists, then the algorithm may be suspected. In that case, compare the effects of different algorithms (e.g., inverse vs. iterative, classical vs. statistical) on the appearance of the artifact. This will verify whether a particular specimen is not restored well by a given algorithm.

If an apparent artifact is visible in both deconvolved and raw images, then some aspect of specimen preparation or optical aberration is implicated. In particular, immunofluorescence staining is often discontinuous along cytoskeletal filaments (actin filaments, microtubules, and intermediate filaments). Several causes can contribute to this problem. One is that cytoskeletal filaments may be masked by cytoskeleton-associated proteins, so antibody accessibility is variable along the filament. Another cause is that the fixatives used in immunocytochemistry often do not faithfully preserve filamentous structures. For example, microtubules may fragment during formaldehyde or methanol fixation. A more faithful fixative such as glutaraldehyde can be tried, but at the expense of antigenicity. Optimal fixation methods are described and discussed in the following sources: (4,14,23,44).

Another type of specimen-dependent artifact derives from RI gradients in the specimen. These cause lensing effects in the specimen before light reaches the objective lens and, therefore, may distort the image. Some specimens (e.g., embryos or thick tissues) have yolk granules or other organelles whose RI is significantly different from their surroundings. In addition, there may be RI gradients in the mounted specimen caused by heterogeneous mixing of the mounting medium. (This is particularly likely if the specimen is mounted in a medium that it was not previously immersed in.) Computational methods for correcting such heterogeneities may be available in the future (33). However, it will always be best to minimize these artifacts by thoroughly immersing the specimen in a mounting medium that is well matched to the specimen's own RI and by adjusting the RI of the immersion medium to compensate for that of the mounting medium.

Horizontal and Vertical Lines

Probably the most common artifacts of deconvolution microscopy are horizontal and vertical lines in x, y, and z. Usually, these lines can be seen in the raw image when examined carefully, meaning that they are not caused by deconvolution but are enhanced by it. However, lines or bands can also be due to edge artifacts at the borders between subvolumes, a form of ringing (discussed above). In any case, such artifacts cannot be mistaken for biological structures and are easily removed.

Horizontal or vertical lines in the x-y plane are often due to "column defects" in the CCD chip used to record the image (31,51). If one pixel in the read register of the chip has a defect or if transfer to that pixel is less efficient, then this will appear as a line perpendicular to the read register. This line is then enhanced by deconvolution. This type of problem can be corrected with a "flat-fielding" correction, also called "shading" or "background" correction (11) and is included in most packages.

Vertical lines seen in x-z or y-z views (called "z-lines") are due to variation in the response characteristics of pixels. Each pixel has a slightly different gain and offset from its neighbors. In extreme cases there are "bad pixels" whose photon response deviates significantly from their neighbors. In such cases, the same pixel systematically deviates from its neighbors throughout a stack of images, leaving a z-line. This problem can be corrected by flat-fielding, or if not, then by special bad pixel routines that search out bad pixels and replace them with the mean of their neighbors (criteria for these are discussed in Reference 11).

In contrast, horizontal lines in x-z or y-z views represent whole planes in the image stack that are uniformly brighter than their neighbors (11). This problem is due to fluctuations in the illumination system: if the arc lamp power output changes during data collection, then this variation will be recorded as a systematic difference in fluorescence intensity between planes. In addition, the arc in a lamp can wander over the surface of the electrodes, causing time-dependent spatial heterogeneities in the illumination of the specimen. These events will be accentuated by deconvolution but can be seen in the raw image. Both problems can be minimized by replacing the arc lamp if it is old.

If lamp fluctuations continue to be a problem even with a new lamp, then a direct measure of arc lamp power fluctuation combined with a polynomial fit of summed pixel intensities plane-by-plane can be used to apply a correction to the raw image (11). Some form of this correction is included in most deconvolution packages. The second problem, changes in the illumination pattern due to arc wandering, is hard to correct by image processing but can be eliminated by installing a fiber-optic scrambler between the arc lamp and the microscope (34).

SUMMARY

Deconvolution algorithms represent a very powerful tool for the biological microscopist. Deblurring algorithms are 2-D methods that remove blur from images. They produce results quickly, but at the expense of signal strength and quantitative accuracy. Image

restoration algorithms are iterative 3-D methods that are somewhat slower but preserve quantitative relationships in image data. Like any method, deconvolution can produce artifacts if not used carefully, but close attention to specimen preparation and image acquisition will usually eliminate them.

APPENDIX: RESOLUTION CRITERIA

Resolution in fluorescence microscopy is often assessed by means of an optical unit called the Rayleigh criterion. This criterion was originally formulated for assessing the resolution of 2-D telescope images, but it has spread into many other areas of optics. It is defined in terms of the minimum resolvable distance between two point sources of light. In a 2-D image, two point sources are resolvable if their Airy disks are distinct. According to the Rayleigh criterion, two Airy disks are distinct if they are farther apart than the distance at which the principal maximum of one Airy disk coincides with the first minimum of the other Airy disk (Figure 5). If the point sources are of equal wavelength, then their Airy disks have the same diameter, and the Rayleigh criterion is then

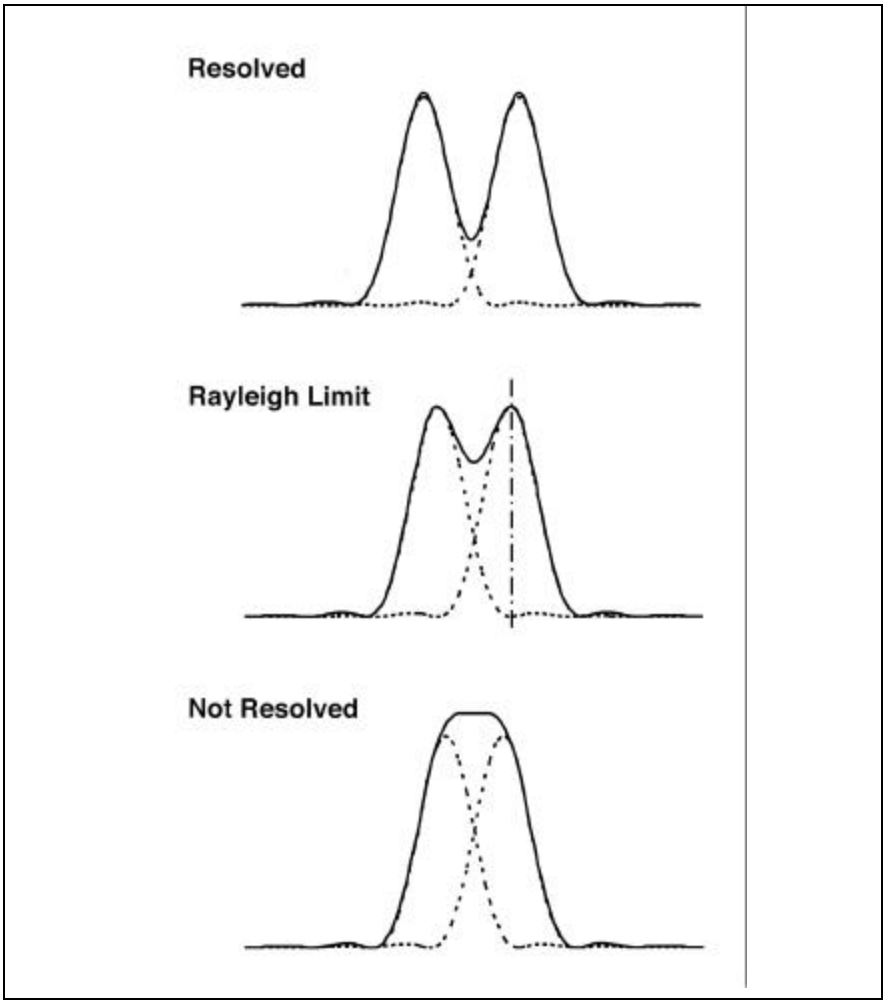

Figure 5. The Rayleigh criterion. Mathematically generated light intensity profiles from two point sources of light. The profiles can be imagined to represent pixel intensities along a line through the Airy disc (i.e., the x-y image of the PSF at focus). The Rayleigh criterion occurs when the maximum intensity of one PSF overlaps with the first minimum of the other. For two points emitting at the same wavelength, this corresponds to the radius (or peak-to-minimum distance) of the central bright region of a single PSF. (Adapted from image by Rod Nave, Department of Physics and Astronomy, Georgia State University, Atlanta, GA, USA. Used by permission.)

equal to the radius of one Airy disk, measured from its point of maximum intensity to its first ring of minimum intensity. For monochromatic images of a given fluorescence wavelength, the Rayleigh criterion can be estimated using a standard formula that is found in many textbooks of optics, including Inoué and Spring (31):

$$\mathbf{d} = 0.61\lambda / \mathbf{NA} \qquad \text{[Eq. 1]}$$

where **d** is the Rayleigh criterion, λ is the wavelength of emitted light, and **NA** is the numerical aperture of the objective lens. Note that the smaller the value of **d**, the higher the resolution.

This formula can be used to assess resolution in the image plane, but not along the optical axis (z-axis). However, an adequate formula for the axial Rayleigh criterion can be deduced using similar reasoning. The minimum resolvable axial distance between two point sources will occur when their axial diffraction patterns are distinct. However, the axial diffraction pattern of a point source is not disk-shaped; rather, it has the hourglass shape or "flare" of the PSF image in x-z or y-z planes. Nonetheless, this hourglass shape has a central bright region, as does the Airy disk. Therefore, to define the axial Rayleigh criterion, we can take the distance from the point of maximum intensity to the first point of minimum intensity of the central bright region along the z-axis. This can be estimated using the following formula:

$$\mathbf{d} = 2\lambda\eta / (\mathbf{NA})^2 \qquad \text{[Eq. 2]}$$

This formula was obtained by comparing the theoretical distribution of light intensity near focus given by Born and Wolf (7) with formulae given by Inoué (30) and Keller (36).

Notice that this formula includes η, the RI of the mounting/immersion media, in the numerator. The mounting and immersion media are assumed to have the same refractive index; otherwise spherical aberration degrades the resolution. (Note that all of these criteria assume aberration-free imaging conditions!) It may be tempting to believe based on this formula that reducing the RI of the immersion medium can improve z-resolution. However, this is a fallacy because the lens NA is also reduced if η is reduced, and since z-resolution varies with the square of NA, the reduction in NA outweighs the reduction in η, and resolution is worse. Notice also that x-y resolution varies only with the first power of NA, whereas z-resolution varies with the square of NA. This means that x-y resolution and z resolution both improve with increasing NA, but z resolution improves more dramatically.

Z-resolution is closely related but not identical to "depth of field". The depth of field is the thickness of the slab of specimen, which appears focused in the image. When looking at a 2-D image, we see a slab of the specimen, with a certain thickness, focused onto a single flat image. Features that seem equally well in focus in the image may reside at different depths in the specimen. The definition of focus is somewhat subjective, but a standard "depth of field unit" is usually defined as half the axial Rayleigh unit (31):

$$\mathbf{d} = \lambda\eta / (\mathbf{NA})^2 \qquad \text{[Eq. 3]}$$

Another criterion that is sometimes used instead of the axial Rayleigh criterion is the full width half maximum (FWHM) of the central bright region of the PSF. Formulae for estimating the FWHM in confocal microscopy are given by Art and Goodman (3). These formulae are identical to those we have produced above for the Rayleigh criterion in wide-field microscopy. We emphasize that these are rough expressions that give practical estimations. They are not exact analytical formulae, which would require vector wave theory.

We would also caution that any resolution criterion is not an absolute indicator of resolution, but rather an arbitrary criterion that is useful for comparing different imaging conditions. The Rayleigh criterion applies specifically to the case when we want to distinguish two self-luminous objects. In other contexts, such as differential interference contrast (DIC), bright-field, or dark-field microscopy, other criteria will apply (31). In some applications such as localization of a moving object, resolution below the Rayleigh limit is possible (e.g., Reference 43). This highlights the fact that resolution is task dependent and cannot be defined arbitrarily for all situations.

In addition, resolution also depends to a great extent on image contrast (i.e., the ability to distinguish signal from background). Think of a picture of London on a foggy day: even with the best high-resolution optics, a gray picture cannot be distinguished from a gray background. Image contrast in the biological context depends mostly on specimen preparation (fixation quality, antibody penetration, evenness of staining, background fluorescence, etc.). Optimizing specimen preparation can improve resolution much more dramatically and at cheaper cost than optics or computers. However, assuming a high-quality preparation, the limit of resolution for any application is always dependent on the PSF, and the Rayleigh criterion gives us at least a basic handle on the size of the PSF.

ACKNOWLEDGMENTS

We are indebted to many colleagues for helpful discussions, including Tim Holmes, Ernst Keller, Steve Paddock, Paul Johannes Helm, Beat Ludin, Michael Wussow, Tom Donnelly, David Carter, Joachim Walter, Mario Moronne, and Chuo-Lung Wang. We also acknowledge the collective assistance of the Confocal Microscopy Listserver (http://listserv.buffalo.edu/archives/confocal.html), and we thank all participants for their generous advice and discussions. W.W. is funded in part by grants from the National Institutes of Health. J.R.S. is a Wellcome Trust Career Development Fellow (054333).

REFERENCES

1. **Agard, D.A.** 1984. Optical sectioning microscopy: cellular architecture in three dimensions. Ann. Rev. Biophys. Bioeng. *13*:191-219.
2. **Agard, D.A., Y. Hiraoka, P. Shaw, and J.W. Sedat.** 1989. Fluorescence microscopy in three dimensions. Methods Cell Biol. *30*:353-377.
3. **Art, J.J. and M.B. Goodman.** 1993. Rapid scanning confocal microscopy. Methods Cell Biol. *38*:62-64.
4. **Bacallao, R., K. Kiai, and L. Jesaitis.** 1995. Guiding principles of specimen preservation for confocal fluorescence microscopy, p. 311-325. *In* J. Pawley (Ed.), Handbook of Biological Confocal Microscopy, 2nd ed. Plenum Press, New York.
5. **Baxes, G.A.** 1994. Digital Image Processing: Principles and Applications. John Wiley and

Sons, New York.
6. **Bohren, C.F.** 1995. Scattering of Particles. *In* M. Bass (Ed.), Handbook of Optics, vol. 1. McGraw-Hill, New York.
7. **Born, M. and E. Wolf.** 1980. Principles of Optics. Pergamon Press, New York.
8. **Carrington, W.** 1990. Image restoration in 3D microscopy with limited data. Proc. Int. Soc. Optical Eng. *1205*:72-83.
9. **Carrington, W.A., R.M. Lynch, E.D. Moore, G. Isenberg, K.E. Fogarty, and F.S. Fay.** 1995. Superresolution three-dimensional images of fluorescence in cells with minimal light exposure. Science *268*:1483-1487.
10. **Castleman, K.R.** 1979. Digital image processing. Prentice-Hall, Englewood Cliffs, NJ.
11. **Chen, H., J.R. Swedlow, M. Grote, J.W. Sedat, and D.A. Agard.** 1995. The collection, processing, and display of digital three-dimensional images of biological specimens, p. 197-210. *In* J. Pawley (Ed.), Handbook of Biological Confocal Microscopy, 2nd ed. Plenum Press, New York.
12. **Conchello, J.A.** 1998. Superresolution and convergence properties of the expectation-maximization algorithm for maximum-likelihood deconvolution of incoherent images. J. Opt. Soc. Am. A. *15*:2609-2619.
13. **Cordeiro, J.M., K.W. Spitzer, W.R. Giles, P.E. Ershler, M.B. Cannell, and J.H. Bridge.** 2001. Location of the initiation site of calcium transients and sparks in rabbit heart Purkinje cells. J. Physiol. *531*:301-314.
14. **Cramer, L. and A. Desai.** Immunofluorescence of the Cytoskeleton. http://iccbweb.med.harvard.edu/mitchisonlab/Pages/gen1.html.
15. **Dernburg, A.F., K.W. Broman, J.C. Fung, W.F. Marshall, J. Philips, D.A. Agard, and J.W. Sedat.** 1996. Perturbation of nuclear architecture by long-distance chromosome interactions. Cell *85*:745-759.
16. **Edwards, A.W.F.** 1972. Likelihood. Cambridge University Press, Cambridge, UK.
17. **Fay, F.S., W. Carrington, and K.E. Fogarty.** 1989. Three-dimensional molecular distribution in single cells analysed using the digital imaging microscope. J. Microsc. *153*:133-149.
18. **Femino, A.M., F.S. Fay, K. Fogarty, and R.H. Singer.** 1998. Visualization of single RNA transcripts in situ. Science *280*:585-590.
19. **Freiden, B.R.** 1975. p. 177-248. *In* T. S. Huang (Ed.), Topics in Applied Physics: Picture Processing and Digital Filtering, vol. 6. Springer-Verlag, New York.
20. **Gibson, S.F. and F. Lanni.** 1991. Experimental test of an analytical model of aberration in an oil-immersion objective lens used in three-dimensional light microscopy. J. Op. Soc. Am. *8*:1601-1613.
21. **Gold, R.** 1964. An Iterative Unfolding Method for Response Matrices. Report no. ANL-6984. Argonne National Laboratory, Chicago.
22. **Goodman, J.W.** 1996. Introduction to Fourier Optics. McGraw-Hill, New York.
23. **Hayat, M.A. (Ed.).** 2000. Principles and Techniques of Electron Microscopy: Biological Applications, 4th ed. Cambridge University Press, Cambridge, UK.
24. **He, X., S. Asthana, and P.K. Sorger.** 2000. Transient sister chromatid separation and elastic deformation of chromosomes during mitosis in budding yeast. Cell *101*:763-775.
25. **Hell, S., G. Reiner, C. Cremer, and E.H.K. Stelzer.** 1993. Aberrations in confocal fluorescence microscopy induced by mismatches in refractive index. J. Microsc. *169*:391-405.
26. **Hiraoka, Y., J.W. Sedat, and D.A. Agard.** 1990. Determination of three-dimensional imaging properties of a light microscope system. Biophysical J. *57*:325-333.
27. **Holmes, T.J.** 1992. Blind deconvolution of quantum-limited imagery: maximum likelihood approach. J. Opt. Soc. Am. A *9*:1052-1061.
28. **Holmes, T.J., S. Bhattacharya, J.A. Cooper, D. Hanzel, V. Krishnamurthi, W.-c. Lin, B. Roysam, D.H. Szarowski, and J.N. Turner.** 1995. Light microscopic images reconstructed by maximum likelihood deconvolution. *In* J. Pawley (Ed.), Handboook of Biological Confocal Microscopy, 2nd ed. Plenum Press, New York.
29. **Holmes, T.J. and N.J. O'Connor.** 2000. Blind deconvolution of 3D transmitted light bright-field micrographs. J. Microsc. *200*:114-127.
30. **Inoué, S.** 1995. Foundations of confocal scanned imaging in light microscopy, p. 1-17. *In* J. Pawley (Ed.), Handbook of Biological Confocal Microscopy, 2nd ed. Plenum Press, New York.
31. **Inoué, S. and K.R. Spring.** 1997. Video Microscopy, 2nd ed. Plenum Press, New York.
32. **Jansson, P.A. (Ed.).** 1997. Deconvolution of Images and Spectra, 2nd ed. Academic Press, New York.
33. **Kam, Z., B. Hanser, M.G.L. Gustafsson, D.A. Agard, and J.W. Sedat.** 2001. Computational adaptive optics for live three-dimensional biological imaging. Proc. Natl. Acad. Sci. USA *98*:3790-3795.
34. **Kam, Z., M.O. Jones, H. Chen, D.A. Agard, and J.W. Sedat.** 1993. Design and construction of an optimal illumination system for quantitative wide-field multi-dimensional microscopy. Bioimaging *1*:71-81.
35. **Karpova, T.S., J.G. McNally, S.L. Moltz, and J.A. Cooper.** 1998. Assembly and function of the actin cytoskeleton of yeast: relationships between cables and patches. J. Cell Biol. *142*:1501-1517.
36. **Keller, H.E.** 1995. Objective lenses for confocal microscopy, p. 111-126. *In* J. Pawley (Ed.), Handbook of Biological Confocal Microscopy, 2nd ed. Plenum Press, New York.
37. **McNally, J.G., T. Karpova, J. Cooper, and J.A. Conchello.** 1999. Three-dimensional imaging by deconvolution microscopy. Methods *19*:373-385.
38. **Oppenheim, A.V. and R.W. Schafer.** 1975. Digital Signal Processing. Prentice-Hall, Englewood Cliffs, NJ.
39. **Oshiro, M.** 1998. Cooled CCD versus intensified cameras for low-light video—applications and relative advantages. Methods Cell Biol. *56*:45-62.
40. **Pawley, J.** 1995. Fundamental limits in confocal microscopy, p. 19-37. *In* J. Pawley (Ed.), Handbook of Biological Confocal Microscopy, 2nd ed. Plenum Press, New York.
41. **Pawley, J.B.** 1994. The sources of noise in three-dimensional data sets, p. 47-94. *In* J. Stevens (Ed.), Three Dimensional Confocal Microscopy. Academic Press, New York.
42. **Preza, C., M.I. Miller, J. Thomas, and J.G. McNally.** 1992. Regularized linear method for reconstruction of three-dimensional microscopic objects from optical sections. J. Opt. Soc. Am. A *9*:219-228.
43. **Qian, H., M.P. Sheetz, and E.L. Elson.** 1991. Single particle tracking. Analysis of diffusion and flow in two-dimensional systems. Biophys. J. *60*:910-921.
44. **Salmon, E.D.** Protocols: Immunofluorescence. http://www.unc.edu/depts/salmlab/protocolsimmunofluorescence.html.
45. **Scalettar, B.A., J.R. Swedlow, J.W. Sedat, and D.A. Agard.** 1996. Dispersion, aberration, and deconvolution in multi-wavelength fluorescence images. J. Microsc. *182*:50-60.
46. **Schaefer, B.C., M.F. Ware, P. Marrack, G.R. Fanger, J.W. Kappler, G.L. Johnson, and C.R. Monks.** 1999. Live cell fluorescence imaging of T cell MEKK2: redistribution and activation in response to antigen stimulation of the T cell receptor. Immunity *11*:411-421.
47. **Schaefer, L.H., D. Schuster, and H. Herz.** Generalized approach for accelerated maximum likelihood based image restoration applied to three-dimensional fluorescence microscopy. J. Microsc. (In Press.)
48. **Shaw, P.J.** 1993. Computer reconstruction in three-dimensional fluorescence microscopy. *In* D. Shotton (Ed.), Electronic Light Microscopy. Wiley-Liss, New York.
49. **Shaw, P.J. and D.J. Rawlins.** 1991. The point-spread of a confocal microscope: its measurement and use in deconvolution of 3-D data. J. Microsc. *163*:151-165.
50. **Shaw, P.J. and D.J. Rawlins.** 1991. Three-dimensional fluorescence microscopy. Prog. Biophys. Mol. Biol. *56*:187-213.
51. **Sluder, G. and D. Wolf (Eds.).** 1998. Methods in Cell Biology, vol. 56. Academic Press, New York.
52. **Straub, M., P. Lodemann, P. Holroyd, R. Jahn, and S.W. Hell.** 2000. Live cell imaging by multifocal multiphoton microscopy. Eur. J. Cell Biol. *79*:726-734.
53. **Swedlow, J.R., J.W. Sedat, and D.A. Agard.** 1997. Deconvolution in optical microscopy, p. 284-309. *In* P.A. Jansson (Ed.), Deconvolution of Images and Spectra, 2nd ed. Academic Press, New York.
54. **Swedlow, J.R., J.W. Sedat, and D.A. Agard.** 1993. Multiple chromosomal populations of topoisomerase II detected in vivo by time-lapse, three-dimensional wide field microscopy. Cell *73*:97-108.
55. **Tikhonov, A.N. and V.Y. Arsenin.** 1977. Solutions to ill-posed problems. Winston, Washington, D.C.
56. **Young, I.T.** 1989. Image fidelity: characterizing the imaging transfer function. Methods Cell Biol. *30*:1-45.

Address correspondence to:

Dr. Jason Swedlow
Division of Gene Regulation and Expression
MSI/WTB Complex
University of Dundee
Dow Street
Dundee DD1 5EH
Scotland
e-mail: j.swedlow@dundee.ac.uk