

into the various class regions. Neural networks can fit very complex functions to describe surfaces separating regions, which may perform better in some cases than simple rectangular prisms, ellipsoids or polyhedra, without encountering the problems of storage, sensitivity to outliers, and slow speed inherent in k NN methods.

Syntactical models

Another approach to feature recognition is used, and it is quite different from the parameter space methods described earlier (Fu, 1974, 1982; Pavlidis, 1977; Tou and Gonzalez, 1981; Schalkoff, 1991). Syntactical models deal primarily with shape information, usually derived from the feature outline or skeleton. This is broken down into the important pieces, and their relative positions recorded. This recording can be likened to a set of characters that spell a word, and compared with other words on the basis of the number of missing, substituted, or rearranged characters.

To give an example, **Figure 34** shows the key elements of the letter A. The two branch points at the sides and the corner or apex at the top identify this letter visually, and no other letter in the alphabet has the same key points nor the same arrangement of them. The “springs” in the figure indicate that the key points can move around quite a bit without changing their usefulness for identification, as long as their relative positions remain in the same topological order. Extracting these key points and their arrangement provides a very robust identifier of the letter, which is quite insensitive to size and font. Some OCR (optical character recognition) programs that convert printed text to computer-readable files use similar methods.

Key points such as branches and ends can be obtained directly from the skeleton, as shown in Chapter 7. Corners can be isolated in several ways. **Figure 35** shows one, in which convolution with a simple 3×3 kernel can select a corner in any specified orientation. A more general corner-finding method traverses the skeleton or outline of the feature using chain code, introduced in Chapter 7. A sequence of links in the chain that gives a net 90-degree (or more) change in direction within a certain (arbitrary) distance along the chain is interpreted as a corner. Fitting functions — polynomials or splines — to the edge points can also be used to locate corners.

Figure 36 shows another example of letters A through E, in several fonts, with the key identifying points of each (ends, branches, and corners). There is actually more information than needed to uniquely identify each character; some of the points can be missing (e.g., due to extreme alterations in letter shape in handwritten characters) without compromising readability. Similar schemes are used in some of the efforts to develop computer programs that read handwriting. Syntactical characterization of the shape and density patterns in human chromosomes is routinely used in karyotyping.




Figure 34. The important branch and corner points that identify the letter A. The springs that connect them indicate that shifting their positions, as long as their relative order remains the same, does not interfere with the identification.

Figure 35. Extracting corner points from a skeleton or outline by using a 3×3 convolution kernel. In general, any pattern can be found by applying a kernel that has the same shape as the desired pattern.





Figure 36. The letters A through E with their key branches, corners, and end points.



3D Image Acquisition

Volume imaging versus sections

Studying the three-dimensional (3D) structures of solid objects is often a goal of imaging. Many of the two-dimensional (2D) images used in the preceding chapters have been sections through three-dimensional structures. This is especially true in the various types of microscopy, where either polished flat planes or cut thin sections are needed in order to form the images in the first place. But the specimens thus sampled are three-dimensional, and the goal of the microscopist is to understand the three-dimensional structure.

There are quantitative tools of great power and utility that can interpret measurements on 2D images in ways that characterize 3D structures. As summarized in Chapter 8, these enable the measurement of volume fraction of phases, surface area of interfaces, mean thickness of membranes, and even the size distribution of spherical particles seen only in a random section. But there are many aspects of structure, both quantitative and qualitative, that are not accessible from 2D section images. Topological properties comprise a major category of such information, including such “simple” properties as the number of separate objects in a volume. Features visible in a single two-dimensional section do not reveal the actual three-dimensional structure present, as shown in **Figure 1**. In Chapter 8, a method was shown for unfolding the distribution of circles to determine the size distribution and total number of spheres in a sample, although this procedure requires making a critical assumption that the objects are all spheres. If the shape of features is unknown, can vary, or is a function of size, then this method fails.

Furthermore, as valuable as numerical measurement data may be, they cannot provide the typical viewer with a real sense of the structure. As pointed out in the first chapter, we are overwhelmingly visual creatures — we need to *see* the 3D structure. Furthermore, our world is 3D. We are accustomed to looking at external surfaces, or occasionally through transparent media, not at thin sections or polished cut surfaces. Try to imagine (again the need to resort to a word with its connotation of vision) standing by a busy street in which an imaginary plane exists transverse to traffic flow, and that you can see that plane but nothing else. What do cars and people look like as they pass through that plane?

If you have a well-developed geometric sense, or experience as a radiologist or draftsman, you may be able to accurately visualize the appearance of that plane as portions of torsos, engine




Figure 1. Examples of sections through three different structures that produce the same 2D sections:

- (a) three discrete objects;
- (b) one object with simple connectivity;
- (c) one object with multiple connectivity.

blocks, even simple shapes such as tires, are cut by the plane. Most people will have difficulty imagining such collections of sections, and when asked to sketch the sections of even simple shapes will produce wildly inaccurate results.

If you doubt this, give yourself a simple test. Get some Cheerios[®], or rotini noodles, or some other simple food object that has a well defined shape (**Figure 2**). Then mix up an opaque matrix (fudge is good) and stir in the objects. While it is hardening, try to draw what various representative random slices through the structure might look like. After it has hardened, cut slices through the sample and compare the actual results to your sketches. My experience with students is that they have a strong tendency to imagine sections through the object that are parallel to principal axes, and ones that pass through the geometrical center of the objects. For the Cheerios (little torii), few actual sections consist of two side-by-side circles or one annulus, and are not necessarily even convex (**Figure 3**). For the rotini, most people do not realize that sections through the curved surfaces actually produce straight lines.

As difficult as this particular visualization task may be, even fewer people can make the transition in the opposite direction: given a collection of section data, to reconstruct in the mind a correct representation of the 3D structure. This is true even of those who feel quite comfortable with the 2D images themselves. Within the world of the 2D images, recognition and understanding can be learned as a separate knowledge base that need not relate to the 3D world. Observing the




Figure 2. Food objects (Cheerios[®] and rotini noodles) that are useful for experimenting with the relationship between 3D shapes and 2D section images. These are somewhat simpler and much more consistent than many natural structures we would like to understand.




Figure 3. A few sections through a torus, producing one or two intersections that may be either convex or concave.

characteristic appearance of dendrites in polished metal samples, or that of mitochondria in thin electron microscope sections of cells, or internal organs in medical images (especially because these latter images are almost always oriented in transaxial, coronal, or sagittal views), does not necessarily mean that the complex 3D shapes of these objects become familiar. In fact, these examples are good illustrations of erroneous 3D interpretations that have persistently been made from 2D section images.

Because of this difficulty in using 2D images to study 3D structure, there is interest in 3D imaging. It may be performed directly, as discussed in this chapter, by actually collecting a 3D set of information all at once, or indirectly by gathering a sequence of 2D (slice) images and then combining them, as discussed in the next chapter. There are a variety of ways to acquire 2D images to assemble the data needed for 3D imaging, and also a great variety of ways to present the information to the user. Many of each are discussed in Chapter 12. The large number of approaches suggests that no one way is best, either for most viewers or for most applications.

Any method that reconstructs internal structural information within an object by mathematically reconstructing it from a series of projections is generally referred to as tomography. It may be used to obtain true 3D arrays of voxels (the 3D analog of the pixel in a 2D image), or to obtain a 2D image from a series of one-dimensional (1D) line projections. The latter method is used in most medical imaging, which is by far the most common application of CT (computed tomography) at the present time, although, the same basic techniques are used for 3D imaging and for other imaging signals.

Medical tomography primarily uses x-ray absorption, magnetic resonance imaging (MRI), positron-emission tomography (PET), and sound waves (ultrasound). Other fields of application and research use many different frequencies of electromagnetic radiation, from x- and gamma rays (nm wavelengths), through visible light and even microwave radiation (wavelengths from cm up). Besides photons, tomography is regularly performed using electrons and neutrons. In addition to absorption of the particles or radiation, tomography can be based on the scattering or emission of radiation as well. Tomography of mixing fluids can sometimes be performed by measuring the electrical resistance or impedance between multiple points around the containment vessel (Mann et al., 2001; York, 2001).

Sound waves produced by small intentional explosions or by “ground thumpers” are used to image underground strata for prospecting, while naturally occurring noise sources such as earthquakes are used to perform seismic tomography, imaging underground faults, rock density beneath volcanoes, and locating discontinuities between the mantle and core of the earth. Devices are also used for listening to moon- and mars-quakes, which will reveal their internal structure, and active studies of the seismic structure of the sun.

Basics of reconstruction

X-ray absorption tomography is one of the oldest, most widely used methods and will be used here to illustrate the various parameters, artefacts, and performance possibilities. Images produced by computerized axial tomography (CAT scans) and similar methods using magnetic resonance, sound waves, isotope emission, x-ray scattering or electron beams, deserve special attention. They are formed by computer processing of information from many individual pieces of projection information obtained nondestructively through the body of an object, which must be unfolded to see the internal structure. The mathematical description of the process presented here is that of x-ray absorption tomography, as it is used both in medical applications and in industrial testing (Herman, 1980; Kak and Slaney, 1988, 2001; Natterer, 2001; Natterer and Wubbeling, 2001). Similar sets of equations and methods of solution apply to the other signal modalities.

Absorption tomography is based on physical processes that reduce intensity as radiation or particles pass through the sample in straight lines. In some other kinds of tomography, the paths are not straight and the reconstruction takes place along curved lines (e.g., magnetic resonance imaging and x-ray scattering tomography), or even along many lines at once (seismic tomography). This makes the equations and graphics slightly more confusing, but does not affect the basic principles involved.

X-rays pass through material but are absorbed along the way according to the composition and density which they encounter. The intensity (number of photons per second) is reduced according to a linear attenuation coefficient μ , which for an interesting specimen is not uniform but has some spatial variation so that we can write $\mu(x,y,z)$, or for a 2D plane through the object, $\mu(x,y)$. The linear attenuation coefficient is the product of the density and the mass absorption coefficient, which depends on the local elemental composition. In medical tomography, the composition varies only slightly, and density variations are primarily responsible for producing images. For industrial applications, significant variations in composition are also usually present. The measured intensity along a straight line path through this distribution is given by

$$\int \mu(x,y) dS = \log_e \frac{I_0}{I_d} \quad (1)$$

where I_0 is the incident intensity (from an x-ray tube or radioisotope) that is known and generally held constant, and I_d is the detected intensity. This is called the ray integral equation and describes the result along one projection through the object.

If a series of parallel lines are measured, either one at a time by scanning the source and detector, or all at once using many detectors, a profile of intensity is obtained, which is called a view. As shown schematically in **Figure 4**, this function is usually plotted as the inverse of the intensity, or the summation of absorption along each of the lines. The function is written as $P(\phi, t)$ to indicate that it varies with position along the direction t as rays sample different portions of the object, and also with angle ϕ as the mechanism is rotated around the object to view it from different directions (or equivalently as the object is rotated).




Figure 4. Illustration of a set of projections through an object at a viewing angle ϕ forming the function P .

Each of the views is a one-dimensional profile of measured attenuation as a function of position, corresponding to a particular angle. The collection of many such views can be presented as a 2D plot or image in which one axis is position t and the other is angle ϕ . This image is called a sinogram or the Radon transform of the 2D slice. **Figures 5** and **6** show a simple example. The construction of a planar figure as shown in **Figure 5a** is called a phantom, and is used to evaluate the important variables and different methods for reconstructing the object slice from the projection information. Such phantoms typically mimic the important structures of real interest, as an aid to evaluating reconstruction algorithms. Compare it to **Figure 5b** showing an actual MRI image of a brain tumor.

The individual projection profiles shown in **Figure 6** show some variation as the angle is changed, but this presentation is difficult to interpret. The sinogram in **Figure 6b** organizes the data so that it can be examined more readily. The name sinogram comes from the sinusoidal variation of position of projections through the various structures within the phantom as a function of rotation, which is evident in the example. The name Radon transform acknowledges the fact that the principles of this method of imaging were published in 1917 by Radon. However, the equations he presented did not provide a practical way to implement a reconstruction since they required a continuous array of projections, and it was not until Hounsfield and Cormack developed a practical reconstruction algorithm and hardware that CAT scans became a routine medical possibility. A. M. Cormack developed a mathematically manageable reconstruction method at Tufts University (Medford, MA) during 1963–1964, and G. N. Hounsfield designed a working instrument at EMI, Ltd. in England in 1972. They shared the Nobel prize in 1979.

The Fourier transform of the set of projection data in one view direction can be written as

$$S(\phi, \omega) = \int P(\phi, t) e^{-j2\pi\omega t} dt \quad (2)$$

Figure 5. A phantom (test object with geometrical shapes of known density) and a real magnetic resonance image of a tumor (bright area) in a section through a head.





Figure 6. Sixteen attenuation profiles for the phantom in Figure 5 (a) and the sinogram or Radon transform (b) produced by plotting 180 such profiles (each as one horizontal line).



Radon showed that this could also be written as

$$S(\phi, \omega) = \iint f(x, y) e^{-j2\pi\omega(x\cos\phi + y\sin\phi)} dx dy \quad (3)$$

which is simply the 2D Fourier transform $F(u, v)$ for the function $f(x, y)$ with the constraints that $u = \omega \cos \phi$ and $v = \omega \sin \phi$. This is consequently the equation of the line for the projection.

What this relationship means is that starting with the original image of the phantom, forming its two-dimensional Fourier transform as discussed in Chapter 5, and then looking at the information in that image along a radial direction from the origin normal to the direction ϕ would give the function S , which is just the 1D Fourier transform of the projection data in direction ϕ in real space. The way this can be used in practice is to measure the projections P in many directions, calculate the one-dimensional transforms S , plot the complex coefficients of S into a 2D transform image in the corresponding direction, and after enough directions have been measured, perform an inverse 2D Fourier transform to recover the spatial domain image for the slice. This permits a reconstruction of the slice image from the projection data, so that a non-destructive internal image can be obtained. It is the principle behind tomographic imaging.

Figure 7 shows an example. Eight views or sets of projections are taken at equal angle steps, the Fourier transform of each is calculated and plotted into a 2D complex image, which is then reconstructed. The image quality is only fair, because of the limited number of views. When 180 views at one-degree intervals are used, the result is quite good. The artefacts which are still present arise because of the gaps in the frequency space image. This missing information is especially evident at high frequencies (far from the origin) where the lines from the individual views become more widely spaced. All tomographic reconstruction procedures are sensitive to the number of views, as we will see.

By collecting enough views and performing this Fourier space reconstruction, it is possible to perform real tomographic imaging. In practice, few systems actually work this way. An exactly equivalent procedure that requires less computation is also available, known as filtered backprojection. This is the method used in most medical scanners and some industrial applications.

The principle behind backprojection is simple to demonstrate. The attenuation plotted in each projection in a view is due to the structure of the sample along the individual lines, or ray integrals. It is not possible to know from one projection just where along the line the attenuation occurs, but it is possible to evenly distribute the measured attenuation along the line. If this is done for only




Figure 7. Reconstruction in frequency space. The (complex) 1D Fourier transforms of projection sets or views at different angles are plotted into a 2D frequency domain image, which is then reconstructed: (a) 8 views in frequency space; (b) 180 views in frequency space; (c) reconstruction from a; (d) reconstruction from b.

a single view, the result is not very interesting. If it is done along projections from several views, however, the superposition of the density or attenuation values should correspond to the features present in the structure.

Figure 8 illustrates this result for the same phantom. It is possible to see the dense (dark) cylinder with its hollow (light) core in the projections. Data from several views overlap to delineate the cylinder in the reconstructed image. There is a problem with this result, however. The attenuation or density of uniform regions in the original phantom is not constant, but increases toward the center of the section. Also, edges are blurred.

The cause of this problem can be described in several different but equivalent ways. The projections from all of the views contribute too much to the center of the image, where all projections overlap. The effect is the same as if the image was viewed through an out-of-focus optical system whose blur or point spread function is proportional to $1/r$, where r is the frequency, or the distance from the center of the frequency transform.

We saw in Chapter 5 how to remove a known blur from an image: apply an inverse function to the frequency space transform that attenuates low frequencies before retransforming. Based on the Fourier approach, and writing the reverse transformation in terms of polar coordinates, this gives


$$f(x, y) = \int_0^{\infty} \int_0^{2\pi} S(\phi, \omega) |\omega| e^{j2\pi\omega(x\cos\phi + y\sin\phi)} d\omega d\phi \quad (4)$$

or, in terms of x and y ,

$$f(x, y) = \int_0^{\pi} Q_{\phi}(x\cos\phi + y\sin\phi) d\phi \quad (5)$$

Figure 8. Back projection in the spatial domain. The attenuation values in each view are projected back through the object space along each projection line. Adding together the data from many view directions does show the major features, but the image is blurred:

- (a) 1 view;
- (b) 6 views;
- (c) 30 views;
- (d) 180 views.



where

$$Q_{\phi} = \int_0^{\pi} S(\phi, \omega) |\omega| e^{j2\pi\omega t} d\omega \quad (6)$$

This is just the convolution of S , the Fourier transform of the projection, by $|\omega|$ the absolute value of frequency. In frequency space, this is an ideal inverse filter which is shaped as shown in **Figure 9**. But, as was pointed out in Chapter 5, convolutions can also be applied in the spatial domain. The inverse transform of this ideal filter is also shown in **Figure 9**. Note its similarity to the shape of a Laplacian or difference of Gaussians as discussed in Chapter 4.

As a 1D kernel or set of weights, this function can be multiplied by the projection P just as kernels were applied to 2D images in Chapters 3 and 4. The weights are multiplied by the values, and the sum is saved as one point in the filtered projection. This is repeated for each line in the projection set or view. **Figure 10** shows the result for the projection data, presented in the form of a sinogram. Edges (high frequencies) are strongly enhanced, and low-frequency information is suppressed.

The filtered data are then projected back, and the blurring is corrected as shown in **Figure 11**. Filtered backprojection using an ideal or inverse filter produces results identical to the inverse Fourier transform method described previously. The practical implementation of filtered back projection is easier because the projection data from each view can be filtered by convolution (a 1D operation) and the data spread back across the image as it is acquired, with no need to store the complex (i.e., real and imaginary values) frequency space image needed for the Fourier method, or retransforming it afterward.

Notice in **Figure 11** that the quality of the image, and the effect of number of views on the artefacts, is identical to that shown for the frequency space method in **Figure 7**. In the absence of noise in the data and other effects which will be discussed in the next section, these two methods are exactly equivalent.





Figure 9. An ideal inverse filter, which selectively removes low frequencies, and its spatial domain equivalent kernel.

Figure 10. The filtered projection data from Figure 6, shown as a sinogram.



Algebraic reconstruction methods

The problem of solving for the density (actually, for the linear attenuation coefficient) of each location in the image can also be viewed as a set of simultaneous equations. Each ray integral (or summation, in the finite case we are dealing with here) provides one equation. The sum of the attenuation coefficients for the pixels (or voxels) along the ray, each multiplied by a weighting factor that takes into account the actual path length of that ray through the pixel, is equal to the measured absorption. **Figure 12** illustrates the relationship between the pixels and the ray integral equations.


The number of unknowns in this set of equations is the number of pixels in the image of the slice through the specimen. The number of equations is the number of ray integrals, which is generally the number of detectors used along each projection profile times the number of view angles. This is a very large number of equations, but fortunately many of the weights are zero (most pixels are not involved in any one particular ray integral equation). Furthermore, the number of equations rarely equals the number of unknowns. But fortunately there are a number of practical and well tested computer methods for solving such sets of sparse equations when they are under- or over-determined.

It is not our purpose here to compare the various solution methods. A suitable understanding of the method can be attained using the simplest of the methods, known as the algebraic reconstruction technique or ART (Gordon, 1974). In this approach, the equations are solved iteratively. The set of equations can be written as

$$A^{m \times n} x^n = b^m \quad (7)$$

Figure 11. Filtered back projection. The method is the same as in Figure 8, except that the values for each view have been filtered by convolution with the function in Figure 9:

- (a) 1 view;
- (b) 2 views;
- (c) 4 views;
- (d) 8 views;
- (e) 16 views;
- (f) 32 views;
- (g) 180 views.



where n is the number of voxels, m is the number of projections, and A is the matrix of weights that correspond to the contribution of each voxel to each ray path. The voxel values are the $\mathbf{x}$ values and the projection measurements are the $\mathbf{b}$ values. The classic ART method calculates each iterative set of $\mathbf{x}$ values from the preceding ones as

$$\mathbf{x}^{k+1} = \mathbf{x}^k + A_i(b_i - A_i^T \mathbf{x}^k) \|A_i\|^{-2} \quad (8)$$

The value of λ , the relaxation coefficient, generally lies between 0 and 2, and controls the speed of convergence. When λ is very small, this becomes equivalent to a conventional least squares solution. Practical considerations, including the order in which the various equations are applied, are dealt with in detail in the literature (Censor, 1983, 1984).

Figure 13 shows a simple example of this approach. The 16×16 array of voxels has been given density values from 0 to 20 as shown in **Figure 13b**, and three projection sets at view angles of





Figure 12. Schematic drawing of pixels (or voxels, because they have depth) in a plane section of the specimen, and the ray integral equations which sum up the attenuation.


Figure 13. Example of the application of an iterative solution. Three projection sets were calculated for an array of 25 detectors, with view directions of 0°, 90°, and 180° (a) The simulated specimen (b) contains a 16 × 16 array of voxels. The calculation results after (c) one, (d) five, and (e) 50 iterations are shown.



0°, 90°, and 180° calculated for the fan beam geometry shown in **Figure 13a**. For an array of 25 detectors, this gives a total of 75 equations in 256 unknowns. Starting with an initial guess of uniform voxels (with density 10), the results after one, five and fifty iterations are shown. The void areas and internal square appear rather quickly, and the definition of boundaries gradually improves. The errors, particularly in the corners of the image where fewer ray equations contain any information, and at the corners of the internal dense square, where the attenuation value changes abruptly, are evident. Still, considering the extent to which the system is underdetermined, the results are rather good.

Kaczmarz' method for this solution is illustrated in **Figure 14**, for the very modest case of three equations and two unknowns, and $\lambda = 1$. Beginning at some initial guess, for instance that all of the pixels have the same attenuation value, one of the equations is applied. This is equivalent to moving perpendicular to the line representing the equation. This new point is then used as a starting point to apply the next equation, and so on. In the real case the equations do not all meet in a perfect point. Because of finite precision in the various measurements, counting statistics,

Figure 14. Schematic diagram of Kaczmarz' method for iteratively solving a set of equations, shown here for the case of two unknowns.



machine variation, etc. there is no single point that represents a stable answer. Instead, the solution converges toward a region that is mostly within the region between the various lines, and then oscillates there. In a high-dimensionality space with some noisy equations, however, it is possible for the solution to leave this region and wander away after many iterations.

In real cases with many dimensions, the convergence may not be very fast. The greatest difficulty in using the iterative algebraic technique is deciding when to stop. Logically, we would like to continue until the answer is as good as it can get, but without knowing the "truth" this stopping point is not possible to judge exactly. Some methods examine the change in the calculated image after each iteration and attempt to judge from that when to stop (for instance, when the normalized total variation in pixel values falls below some arbitrary limit, or when it increases from the previous iteration). This method is prone to serious errors in a few cases, but is used nonetheless. It should be noted that the penalty for continuing the iteration is not simply the computational cost, but also the possibility that for some sets of data, the answer may start to diverge (leave the bounded region near the crossover point). This condition is, of course, highly undesirable.

Given the drawbacks to the algebraic approach, and the relative simplicity and straightforward approach of the filtered back projection method, why would we use this method? There are several potential advantages of algebraic methods such as ART. First, the filtered back projection method, and the Fourier transform method which it embodies, require that the number of views be rather large, and that they be equally spaced so that frequency space is well filled with data. Missing angles, or entire sets of angles which may be unattainable due to physical limitations, present problems for filtered back projection and introduce significant artefacts. ART methods can still produce an acceptable reconstruction. There may still be lack of detail in portions of the reconstructed image which are undersampled by the projections, but the artefacts do not spread throughout the entire image. In fact, acceptable reconstructions are often obtained with only a very few views (examples are shown in the following).

Another advantage to ART is the ability to apply constraints. For instance, it is possible in a filtered back projection or Fourier transform method to calculate negative values of density (attenuation) for some voxels, because of the finite measurement precision. Such values have no physical meaning. In the iterative algebraic method, any such values can be restricted to zero. In the schematic diagram of **Figure 14**, this amounts to restricting the solution to the quadrant of the graph with positive values.

In fact, any other prior knowledge can also be applied. If it is known that the only possible values of density and attenuation in the specimen correspond to specific materials, then the values can be easily constrained to correspond. Any geometric information, such as the outside dimensions of the object, can also be included (in this case, by forcing the voxels outside the object boundaries to zero).

It is even possible to set up a grid of voxels that are not all of the same size and spacing. This setup might allow, for instance, the use of a fine voxel spacing in the interior of an object where great detail is desired, but a much coarser grid outside (or vice versa). This would still allow the calculation of the contribution of the outside material to the ray integrals, but would reduce the number of unknowns to produce a better solution for any given number of views and projections. Sets of non-square pixels or non-cubic voxels can also be used when necessary to conform to specific object shapes and symmetries.

The flexibility of the algebraic method and its particular abilities to use *a priori* information often available in an industrial tomography setting compensates for its slowness and requirements for large amounts of computation. The calculation of voxel weights (the **A** matrix) can be tedious, especially for fan beam or other complex geometries, but no more so than backprojection in such cases, and it is a one-time calculation. The use of other solution methods than the simple iterative approach described here can provide improved stability and convergence.

Maximum entropy

There are other ways to solve these huge sets of sparse equations. One is the so-called maximum entropy approach. Maximum entropy was mentioned as an image processing tool to remove noise from a 2D image, in Chapter 3. Bayes' theorem is the cornerstone for the maximum entropy approach, given that we have relevant prior information which can be used as constraints. In the case where no prior information is available but noise is a dominant factor, Bayes' theorem leads to the classical or "least squares" approximation method. It is the use of prior information that permits a different approach.

The philosophical justification for the maximum entropy approach comes from Bayesian statistics and information theory. It has also been derived from Gibbs' concept of statistical thermodynamics (Jaynes, 1967). For the nonspecialist, it can be described as follows: find the result (distribution of brightnesses in pixels of image, distribution of density values in a voxel array, or practically anything else) that is feasible (consistent with the known constraints, such as the total number of photons, the non-negativity of brightness or density at any point, the physics involved in the detector or measurement process, etc.) and has the configuration of values which is most probable.

This probability is defined as being able to be formed in the most ways. For an image formed by photons, all photons are considered indistinguishable, and the order in which they arrive is unimportant, so the distribution of photons to the various pixels can be carried out in many ways. For some brightness patterns (images) the number of ways to form the pattern is much greater than others. We say that these images with greater multiplicity have a higher entropy. Nature can form them in more ways, so they are more likely. The entropy is defined as $S = -p_i \log p_i$, where p_i is the fraction of pixels with brightness value i .

The most likely image (from a simple statistical point of view) is for all of the pixels to get the same average number of photons, producing a uniform grey scene; however, this result may not be permitted by our constraints, one of which is the measured brightness pattern actually recorded. The difference between the calculated scene and the measured one can only be allowed to have a set upper limit, usually based on the estimated noise characteristics of the detector, the number of photons, etc. Finding the feasible scene which has the highest multiplicity is the maximum entropy method.

For instance, in solving for the tomographic reconstruction of an object from the set of ray integral equations obtained from various view angles, we have a large set of simultaneous equations in many unknowns. Instead of formally solving the set of simultaneous equations, for instance by a traditional Gauss-Jordan elimination scheme which would take far too many steps to be practical,

the maximum entropy approach recasts the problem. Start with any initial guess (in most "well-behaved" cases, the quality of that guess matters little in the end result) and then iteratively, starting at that point, find another solution (within the class of feasible solutions as defined by the constraints) that has a higher entropy. Deciding which way to move in the space defined by the parameters (the values of all the voxels) is usually done with LaGrange multipliers by taking partial derivatives and trying always to move "uphill" where the objective function used to evaluate each set of values is the entropy.

It is usually found that the solution having the maximum feasible entropy (i.e., permitted by the constraints) is hard against the boundary formed by those constraints, and if they were relaxed the solution would move higher (towards a more uniform image). Knowing or assuming that the solution lies along the constraint boundaries allows use of more efficient schemes for finding the best solution. For the noise removal problem discussed in Chapter 3, the constraint is commonly the chi-squared value of the smoothed image as compared to the measured one. This is generally assumed to be due to classical noise, and so should have an upper limit and a known distribution.

For tomographic reconstruction, the constraints are based on satisfying the ray integral equations. These are not all consistent, so a weighting scheme must be imposed on the error; linear weighting is the simplest and most often used. It turns out that in most cases, the cluster of solutions with high entropies, all permitted by the constraints, are virtually indistinguishable. In other words, the maximum entropy method does lead to a useful and robust solution. While the solution is still iterative, the method is quite efficient as compared to other solution techniques.

Defects in reconstructed images

The reconstructed example shown above in **Figures 7** and **11** was calculated using projection data simulated by computation, with no noise or any other defects. In real tomography, a variety of defects may be present in the projection sets that propagate errors back into the reconstructed image. Using the same phantom, several of the more common ones can be demonstrated.

Ideally, a large number of view angles and enough detector positions along each projection set will be used to provide enough information for the reconstruction. In the event that fewer projections in a set or fewer views are used, the image has more reconstruction artefacts and poorer resolution, definition of boundaries, and precision and uniformity of voxel values. **Figure 15** shows the effect of fewer projections in each set but still uses 180 view angles. The reconstructed images are displayed with 100×100 pixels. This ideally requires a number of ray integrals in each projection set equal to at least $\sqrt{2}$ times the width, or 141 equations for each view. With fewer, the resolution of the reconstruction degrades.

If fewer view angles are used (but the angular spacing is still uniform), the artefacts in the reconstruction increase as was shown above in **Figure 11**. If the view angles are not uniformly spaced, the results are much worse as shown in **Figure 16**.

In real images, the number of x-ray photons detected at each point in the projection set is subject to fluctuations due to counting statistics. In many cases, both in medical and industrial tomography, the number of photons is limited. In medical applications, it is important to limit the total exposure to the subject. In industrial applications, the limitation is due to the finite source strength of either the x-ray tube or radioisotope source, and the need to acquire as many views as possible within a reasonable time. In either case, the variation in the number of detected x-rays varies in a Gaussian or normal distribution whose standard deviation is the square root of the number counted. Counting an average of 100 x-rays produces a variation whose standard deviation is 10% ($\sqrt{100} = 10$), while an average of 10,000 x-rays is needed to reduce the variation to 1% ($\sqrt{10^4} = 10^2$).





Figure 15. Effect of number of ray integrals in the projection set on reconstructed image quality. Each image is reconstructed with 100×100 pixels, and is calculated from 180 view angles. The images show the use of (a) 25; (b) 49; (c) 75, and (d) 99 ray projections, respectively.

Figure 16. Effect of using a set of view angles that do not uniformly fill the angular range: (a) 150-degree coverage; (b) 120-degree coverage; (c) 90-degree coverage; (d) a different 90-degree range.



The process of reconstruction amplifies the effect of noise in the projections. This is the same effect of noise seen in Chapter 5 for removing blur by deconvolution, and arises from the same mathematical procedures. The filtering process suppresses the low frequencies and keeps the high frequencies, and the counting fluctuations vary randomly from point to point and so are represented in the highest frequency data. **Figure 17** shows the result. Adding a statistical or counting fluctuation of a few percent to the simulated projection data produces a much greater noise in the reconstructed image. Although the density differences in the three regions of the phantom vary by 100%, some of the regions disappear altogether when 10% or 20% noise is added to the projection data.


Suppression of the high-frequency noise in the projection data by the filtering process can reduce the effect of the noise somewhat, as shown in **Figure 18**. Notice that the noise variations in the reconstructed images are reduced, but that the high frequency data needed to produce sharp edges and reveal the smaller structures are gone as well.

Several different filter shapes are used for this purpose. **Figure 19** shows representative examples, in comparison to the shape of the ideal inverse filter that was discussed above. The plots are in terms of frequency. All of the filters reduce the low frequency values, which is required in order to prevent blurring, and all of the noise reduction filters also attenuate the high frequencies in order to suppress the noise.

Another important source of errors in the reconstruction of images is imprecise knowledge of the location of the center of rotation, or variation in that center due to imperfect mechanical mechanisms (Barnes et al., 1990). As shown in **Figure 20**, this variation also produces an effect in the reconstructed image which is magnified. The characteristic "U" shaped arcs result from an off-center rotation, because view angles in a range of 180° were used. If 360-degree rotation is used, a complete circular arc is present (**Figure 21**) that also distorts the reconstruction but is more difficult to recognize. Note that it is not common to collect data over a complete 360-degree set of

Figure 17. Effect of counting statistics on reconstruction. The images were reconstructed from simulated projection data to which Gaussian random fluctuations were added:

- (a) 2%;
- (b) 5%;
- (c) 10%;
- (d) 20%.






Figure 18. Reconstructions from the same projection data with superimposed counting statistics variations as in **Figure 17**, but using a Hann filter instead of an ideal inverse filter to reduce the high-frequency noise.

angles, because in the absence of off-center rotation or beam hardening (discussed below) the second half of the data would be redundant. The effect of a variable center is equal in magnitude with 360-degree rotation, but harder to recognize. In general, it is required that the location of the center of rotation and its constancy should be less than about one-tenth of the expected spatial resolution in the reconstructed images.

Similarly, the motion of the object should be restricted during the collection of the multiple views. In medical imaging, where the most common imaging geometry is to place the subject inside a circular track on which the source and/or detectors rotate, this means that the patient must not breathe during the scanning process. If the images are expected to show the heart, the entire collection of views must be obtained in a time much shorter than the single beat, or else some stroboscopic triggering scheme must be used to collect images at the same relative timing over many beats.

Figure 19. Filter profiles for noise reduction in filtered back projection:

Ideal inverse: $Weight = |f|$

Hann: $Weight = |f| \cdot \{0.5 + 0.5 \cos[(\pi/2)(f/f_m)]\}$

Hamming: $Weight = |f| \cdot \{0.54 + 0.46 \cos(\pi f/f_m)\}$

Butterworth ($n = 3$): $Weight = |f| \cdot 1 / (1 + (f/2f_m)^{2n})$




Figure 20. Effect of errors in the center of rotation on the reconstructed image. In each of these images, the center is consistent, but displaced from the location assumed in the rotation by a fraction of the image width:

- (a) 0.5%;
- (b) 1.0%;
- (c) 2.5%;
- (d) 5.0%





Figure 21. Repeating the reconstructions of Figure 20 using the same number of views (180) but spread over 360° instead of 180°, with the center of rotation displaced from the assumed location by a fraction of the image width:

- (a) 0.5%;
- (b) 1.0%;
- (c) 2.5%;
- (d) 5.0%

Beam hardening

Beam hardening is the name used to describe the effect in which the lower energy or “softer” x-rays from a polychromatic source such as a conventional x-ray tube are preferentially absorbed in a sample. The consequence is that the effective attenuation coefficient of a voxel is different depending on whether it is on the side of the specimen near the source or farther away. This variation along the path is indicated schematically in **Figure 22**. Beam hardening is not a major problem in medical tomography because the variation in composition of the various parts of the human body is only slight. Everything is mostly water with some addition of carbon, a trace of other elements, and for bones, some calcium. The density is variable, and is in fact what the reconstructed image shows, but the range of variation is small. This uniformity makes x-ray tubes an acceptable source, and simple back projection a suitable reconstruction method.

Industrial applications commonly encounter samples with a much greater variation in composition, ranging across the entire periodic table and with physical densities that vary from zero (voids) to more than ten times the density of biological tissue. This large range of variation makes beam hardening an important problem. One solution is to use a monochromatic source such as a radioisotope or a filtered x-ray tube. Another is to use two different energies (Schneberk et al., 1991) or a combination of absorption and x-ray scattering data (Prettyman et al., 1991) and to use the two projection sets to correct for the change in composition in the reconstruction process; however, this method increases the complexity significantly and requires an algebraic solution method rather than back projection or Fourier techniques.

Figure 23 shows a representative example of the beam hardening effect in the same phantom used previously. In this case, the sample composition is specified as void (the lightest region and the surroundings), titanium (the medium grey region of the elliptical object), and iron (the dark region). The total width is 1 cm, and the x-ray tube is assumed to be operating at 100 kV. This is in fact a very modest amount of beam hardening. A larger specimen, a lower tube voltage, higher atomic number elements, or a greater variation in atomic number of density, would produce a much greater effect.

Figure 24 shows reconstructions of the image using view angles which cover 180° and 360° , respectively. In most tomography, 180° is adequate since the projections are expected to be the same regardless of direction along a ray path. This assumption is not true in the case of beam hardening (as it was not for the case of off-center rotation), and so better results are obtained with a full 360° of data. Notice, however, that artefacts are still present. This is particularly true of the central feature, in which the narrow void is hardly visible. **Figure 25** shows the same phantom with no beam hardening, produced with a monochromatic x-ray source.




Figure 22. Schematic diagram of beam hardening. The energy spectrum of x-rays from an x-ray tube is shown at the beginning, middle, and end of the path through the specimen. As the lower energy x-rays are absorbed, the attenuation coefficient of the sample changes independent of any actual change in composition or density.

Figure 23. Example of beam hardening effect on the sinogram or Radon transform

(a) and the inverse filtered data **(b)**. Notice that the contrast of each feature changes according to where it lies within the rotated object.





Figure 24. Reconstruction of the beam-hardened data from Figure 23:

(a) 180 views covering 180°;
(b) 180 views covering 360°.

When x-rays pass through material, the attenuation coefficient that reduces the transmitted intensity consists of two principal parts: the absorption of the x-rays by the excitation of a bound electron, and the scattering of the x-rays either coherently or incoherently into a different direction. In either case, the photons are lost from the direct ray path, and the measured intensity decreases. In the case of scattering, however, the x-rays may be redirected to another location in the detector array (see the discussion of the geometries of various generations of instrument designs, in the following paragraphs).

When this scattering takes place, the measured projection profiles contain additional background on which the attenuation data are superimposed. The presence of the background also produces artefacts in the reconstruction as shown in **Figure 26**. The effect is visually similar to that produced by beam hardening.

In addition, the uniform regions in the object are reconstructed with a variable density due to the background. **Figure 27** shows this reconstruction for a simple annular object (a crude model for a bone cross-section), and **Figure 28** shows plots across the center of the reconstructions. The deviation from a uniform density in the reconstruction is called cupping. Note that this example uses





Figure 25. Reconstruction of the same phantom as **Figure 23** but using a monochromatic 50 kV x-ray source. Notice particularly the void in the center of the object, which is not visible in **Figure 24**.

Figure 26. Reconstruction of the phantom in **Figure 25** when the measured projection sets include scattered background radiation of (a) 5; (b) 10; (c) 20; and (d) 40% of the average intensity. The effect on the image is similar to beam hardening. Small features are obscured by artefacts, and the overall contrast changes.



materials similar to those in the human body. Medical tomography is not usually required to produce a quantitatively accurate measure of density; however, it is used to show the location of internal structures and boundaries (this may not be true in all cases; for example radiography and computed tomography are used to measure bone density loss due to osteoporosis). Industrial tomography is often called upon to measure densities accurately so as to quantify gradients in parts due to processing, and this source of error is therefore of concern.

Although medical applications rarely need to measure densities exactly, they do require the ability to show small variations in density. A test phantom often used to demonstrate and evaluate performance in this category is the Shepp and Logan (1974) head phantom. Composed of ellipses with densities close to 1.0, it mimics in simplified form the human head, surrounded by a much denser skull, and containing regions very slightly lower or higher in density that model the brain structure and the presence of tumors. The ability to image these areas is critical to the detection of anomalies in real head scans.

Figure 27. Reconstruction of a simple annulus with outer composition of calcium carbonate (an approximation to bone) and an inner composition of water (an approximation to tissue): (a) no scattered background; (b) 10% scattered background in projection data.





Figure 28. Line profiles of the density in the images in Figure 27.

Figure 29. Shepp and Logan phantom, intended to represent the difficulty of visualizing a tumor inside the human head. The regions of varying density inside the "brain" range of relative densities of 1.0 to 1.04, while the "skull" has a density of 2.0. They are not visible in the reconstructed image (a) unless some contrast expansion is applied. Here, histogram equalization is used (b) to spread the grey scale nonlinearly to show the various ellipses and their overlaps (and also to increase the visibility of artefacts in the reconstruction). The brightness histograms (c, d) show the effect of the equalization.




Figure 29 shows a reconstruction of this phantom. Using the full dynamic range of the display (values from 0 to 255) linearly to represent the image does not reveal the internal detail within the phantom. Applying histogram equalization (as discussed in Chapter 4) expands the contrast in the center of the histogram so that the different regions become visible. The figure shows the cumulative histograms of display brightness for the original and histogram-equalized images. In the former, the steps show the principal density values; after equalization the values approximate a straight line, and the steps show the distinction between similar regions. A profile plot across the center of the structure shows the different regions with quite uniform density values (**Figure 30**).

As noted previously, tomography can be performed using modalities other than x-ray absorption. One is emission tomography, in which a radioactive isotope is placed inside the object and then reveals its location by emitting gamma ray photons. Collimated detectors around the object can specify the lines along which the source of the photons lie, producing data functionally equivalent to the attenuation profiles of the conventional case.

Figure 31 shows an example of emission tomography using real data, in which another artefact is evident. The bright areas in the reconstruction are cavities within a machined part that contain a radioactive isotope. The sinogram shows the detected emission profiles as a function of view angle. Notice that the width of the regions varies with angle. This variation is due to the finite width of the collimators on the detectors, which cover a wider dimension on the far side of the object as indicated schematically in **Figure 32**. This effect is also present in x-ray absorption tomography, due to the finite size of apertures on the source and the detectors. If the angle of the collimators is known, this effect can be included in the reconstruction, either by progressively spreading the data as the filtered profile is spread back across the voxel array, or by adjusting the voxel weights in the algebraic reconstruction technique.

Figure 30. Brightness profiles across the images in **Figure 29**, showing the uniformity and sharpness of transitions for the regions, and the effect of histogram equalization.






Figure 31. Emission tomography. The sample is a block of aluminum containing several cylindrical cavities containing radioactive cobalt. The detector collects a series of intensity profiles as a function of rotation, shown in the form of a sinogram (**a**). Note that the width of the trace for each cylinder varies as the sample is rotated, due to the finite angle of the entrance collimator to the detector. The reconstruction of the cross-section is shown in (**b**).

Figure 32. Schematic diagram showing the effect of a finite collimator angle on the dimensions and voxels covered in different parts of the object.



Imaging geometries

First-generation tomographic systems collected projection sets at a variety of view angles by moving a source and detector just as shown above in **Figure 3**. **Figure 33** shows the procedure used to collect a complete set of projection data. This is called a pencil-beam or parallel-beam geometry, in which each ray integral is parallel and the projection set can be directly backprojected. It is not very efficient, since only a small solid angle of the generated x-rays can be used, and only a single detector is in use, but is still used in a few industrial imaging situations where time of data acquisition is not a major concern.

Second-generation instruments added a set of detectors so that a fan beam of x-rays could be detected and attenuation measured along several lines at the same time, as shown in **Figure 34**. This procedure requires fewer view angles to collect the same amount of data, but the attenuation measurements from each detector are actually for different angles and there is some re-ordering of the data needed before it can be used.

The so-called fan-beam geometry is appealing in its efficiency, and the next logical step, in so-called third-generation instruments used for medical imaging, was to use a larger array of detectors (and to arrange them on an arc so that each covered the same solid angle and had normal x-ray incidence) and a single x-ray tube. The detectors and tube rotate together about the object as the x-ray tube is pulsed to produce the series of views (**Figure 35**). In fourth-generation systems, a complete ring of detectors is installed and only the source rotates (**Figure 36**). Notice that the x-rays are





Figure 34. Second-generation geometry. The detector array simultaneously measures attenuations in a fan beam, requiring fewer view angles than first generation systems to collect the data.

Figure 33. First-generation geometry. The detector and source move together to collect each projection set, and rotate to many view angles to collect all of the required data.






Figure 35. Third-generation geometry. The x-ray tube and detector array rotate together around the object being imaged, as the tube is rapidly pulsed to produce each view.

no longer normally incident on the detectors in this case. There is a fifth-generation design in which even less hardware motion is required: the x-rays are generated by magnetically deflecting an electron beam against a fixed target ring, rotating the source of x-rays to produce the same effective geometry as in fourth-generation systems, but with even shorter exposure times.

These latter types of geometry are less used in industrial tomography, since they are primarily intended for imaging speed, to minimize exposure and acquire all of the projections before anything can move in the person being imaged. First- or second-generation (pencil or fan beam) methods in which a series of discrete views are collected provide greater flexibility in dealing with industrial problems. All of the methods are equivalent, however, if the various ray integrals using individual detectors in the fan-beam geometry are sorted out according to angle, and either backprojected, used in a Fourier transform method, or used to calculate an algebraic reconstruction with appropriate weights.

Some imaging technologies use different geometries and reconstructions. For instance, ultrasound images, most familiarly used for examining a fetus in utero, are obtained as fan-beam sections, but can be reconstructed to show surfaces (**Figure 37**). Positron-emission tomography (PET), single-photon

Figure 36. Fourth-generation geometry. The detector array forms a complete ring and is fixed. The x-ray tube rotates around the object and is pulsed. Data from the detectors is sorted out to produce the projection sets.




Figure 37. Surface reconstruction from ultrasound scans (image courtesy William Beaumont Hospital).

emission spectroscopy (SPECT), and magnetic resonance imaging (MRI) are similar in geometry to CT scans, but record different information about internal structure.

PET and SPECT localize specific atomic isotopes by their decay. In PET, the positron produced by the atomic nucleus gives rise to a pair of 511 keV photons that travel in exactly opposite directions. Detecting them with a ring of detectors locates the original isotope tag somewhere on the line between the two points, which becomes the line integral used in the reconstruction. SPECT detects gamma rays from isotope decay with collimated detectors. An array of detectors collects information from a parallel set of projections, and is scanned across and around the subject to collect multiple projections.

When protons are placed in a magnetic field, they oscillate (resonate — thus “magnetic resonance imaging” or MRI) at a frequency that depends on the field strength, and absorb energy at the oscillation frequency. This energy is reradiated as the protons return to their ground state. The reradiation involves processes (relaxation of the magnetization components parallel and perpendicular to the field) with different time constants T_1 and T_2 . MRI signal strength depends on the proton concentration (essentially the water concentration in the tissue for medical imaging) but the contrast depends on T_1 and T_2 , which are strongly influenced by the fluid viscosity or tissue rigidity. Weighting the combination of the two signals provides control over the observed image.

PET and MRI data can be very complementary to the density information from x-ray tomography, and in many cases scans using the different modalities are combined by overlaying the section images in registration (**Figure 38**).

Three-dimensional tomography

Although the most common application of tomography is still to form images of planar sections through objects without physical sectioning, the method can be directly extended to generate complete three-dimensional images. Chapter 10 shows several examples of 3D displays of volume data. Most of these, including many of the tomographic images, are actually serial section images.




Figure 38. Co-registration of PET and MRI scans (images courtesy K. A. Paller, Northwestern University).

Whether formed by physical sectioning, optical sectioning (for instance, using the confocal light microscope), or conventional tomographic reconstruction, these methods are not true 3D data sets.

The distinction is that the pixels in each image plane are square, but as they are extended into the third dimension as voxels, they do not necessarily become cubes. The distance between the planes, or the depth resolution, is not inherently the same as the pixel size or resolution within the plane. In fact, few of these methods have depth resolution that is even close to the lateral resolution. Some techniques such as physical or optical sectioning have much poorer depth resolution. Others such as the secondary ion mass spectrometer have depth resolution that is far better than the lateral resolution of images. This has profound effects for 3D image presentation, for image processing, and especially for three-dimensional structural measurement.

True 3D imaging is possible with tomographic reconstruction. The object is represented by a three-dimensional array of cubic voxels, and the individual projection sets become two-dimensional arrays (projection images). Each projection is from a point, and is referred to as a cone-beam geometry (by analogy to the fan-beam method used for single slice projections). The set of view directions must include orientations that move out of the plane and into three dimensions, described by two polar angles. This does not necessarily require rotating the object with two different polar angles, since using a cone-beam imaging geometry provides different angles for the projection lines, just as a fan beam geometry does in two dimensions. The best reconstructions, however, are obtained with a series of view angles that cover the 3D orientations as uniformly as possible.

Several geometries are possible. One of the simplest is to rotate the sample about a single axis as shown in **Figure 39**. This method offers the advantage of precise rotation, since as seen before, the quality of the reconstruction depends on the consistency of the center of rotation. On the other hand, artefacts in the reconstructed voxels can be significant, especially in the direction parallel to the axis and near the north and south poles of the sample. The single-axis rotation method is most often used with x-ray, neutron, or gamma ray tomography because the samples may be rather large and are relatively equiaxed so that the distance that the radiation must pass through the sample is the same in each direction. Improved resolution in the axial direction can be obtained using a helical scan (Wang et al., 1991) in which the specimen rotates while moving in the axial direction (**Figure 40**). This has become the preferred geometry, particularly for small industrial objects.

For electron tomography, most samples are thin sections, and few transmission electron microscope stages permit complete movement of the sample about its axis. For a sample that is essentially




Figure 39. Geometry for volume imaging using radial projections obtained by rotating the sample about a single axis.

Figure 40. In helical scanning, the specimen is raised as it rotates so that multiple views measure absorption along different angles through the voxels.




Figure 41. Tilting the sample in a conical pattern produces a series of projections used for reconstruction in transmission electron microscopy. The spacing along the cone trace is generally not uniform.




slab-like, the geometry that is usually adopted is a series of tilt angles that project along the limbs of a cone, as shown in **Figure 41**. Collecting these projections by controlling the tilt and rotation of the sample in its holder with enough precision to allow good reconstructions is very difficult. Some TEM samples consist of many repeating structures (macromolecules, virus particles, etc.), and a single orientation of the sample can collect enough different views to be used for reconstruction. Because of the use of many different individual objects with various orientations, this method is called random-conical, as compared with the use of equal angular increments. The very small aperture angle of the beam in the TEM produces essentially parallel rather than cone-beam projections, which does simplify the reconstruction and makes backprojection straightforward, but the use of a limited set of views arranged in a cone produces artefacts because little information is available in the axial direction (sometimes referred to as the missing cone of information). Frank (1992) presents a thorough review of the current state of the art in electron microscope tomography.

From a theoretical viewpoint, the best reconstruction for any given number of projections is obtained when they are uniformly distributed in 3D space (**Figure 42**). Constructing a mechanism to achieve accurate rotations about two precisely centered axes is difficult and this technique not widely used.

The reconstruction can be performed with any of the methods used in two dimensions. For Fourier inversion, the frequency space is also a 3D array, and the 2D images produced by each projection are transformed and the complex values plotted on planes in the array. As for the 2D case, filling the space as completely and uniformly as possible is desirable. The Fourier inversion is performed in three dimensions, but this is a direct extension of methods in lower dimensions and in fact the inversion can be performed in one dimension at a time (successively along rows in the u , v , and w directions).

Backprojection can also be used for 3D reconstruction, and as in the 2D case is simply an implementation of the Fourier transform mathematics. The filtering of the 2D images must be performed

Figure 42. Optimum 3D reconstruction is possible when a series of 3D projections is used, by rotating the sample about two axes.



with a 2D convolution, which can be carried out either by kernel operation in the spatial domain or by multiplication in the Fourier domain. The principal difficulty with the backprojection method is that calculation of the matrix of weights can be quite tricky for cone-beam geometry, especially when combined with helical scanning. These values represent the attenuation path length along each of the ray integrals through each of the voxels. The use of backprojection requires a large number of views to avoid artefacts and is most commonly used with single-axis rotation or with helical scans about a single rotational axis, with either cone-beam or parallel-beam projections (Feldkamp et al., 1984; Smith, 1990; Shih et al., 2001). It is difficult to apply to a full 3D set of cone beams because they are spaced at relatively large angles.

Algebraic reconstruction technology (ART) methods are also applicable to voxel arrays. The difficulty in obtaining a uniform set of view angles, which is particularly the case for electron microscopy, can make ART methods more attractive than the inverse Fourier approach. In fact, when using an iterative technique such as ART, the best results are often obtained with a surprisingly small number of views. **Figures 43** and **44** show an example. The specimen (about 2 cm on a side) consists of three different metal cylinders in a plastic block. Chromium, manganese and iron are consecutive elements in the periodic table, with similar densities. Tomographic reconstruction from only 12 views with 3D rotations, using a low power industrial x-ray source, shows the inserts quite well (Ham, 1993).

Of course, more views should produce a better reconstruction. But in most tomography situations, the total dose is a fixed constraint. In some cases, this can be because of concerns about radiation

Figure 43. Geometry of a test sample. Three cylindrical inserts of different metals (Cr, Mn, and Fe) are placed in a plastic block (2 cm on a side).





Figure 44. Tomographic reconstruction of three planes in the xy , yz , and zx orientations passing through the metal inserts in the plastic block shown in **Figure 43**, reconstructed using just 12 cone beam projections with rotations in 3D.

damage to the sample. Dosage to the sample is a concern for medical x-ray tomography, of course; but it also creates problems for electron tomography. The amount of energy deposited in each cubic nanometer of the sample from a focused electron beam is great enough to cook biological tissue, disrupt molecules, and change the structure we want to image.


Even for industrial tomography, the total flux of radiation that can be generated and the time spent acquiring the images usually is limited. A necessary trade-off must be made between the number of projections and the time spent acquiring each one. More time on each projection improves the statistical quality of the view image, so acquiring more projections makes each one noisier, and vice versa. In some experiments with a limited total photon budget, the best quality reconstructions with full 3D rotation were obtained with only 12 projections (Ham, 1993). This small number of views requires an iterative method rather than backprojection.

The limited number of photons becomes particularly critical when low-intensity sources are used. Synchrotrons are excellent sources of x-rays with high brightness and the ability to select a specific monochromatic energy, but are not usually conveniently available for tomographic work. Radioactive sources of gamma rays present handling difficulties and have low intensities as well. X-ray tubes are a convenient source for tomography, with adjustable voltage and a variety of target materials that emit different x-ray spectra. Such a source is not monochromatic, which may cause significant beam hardening effects for many specimens, as discussed previously.

Absorption filters can be used to select just a single band of energies from a polychromatic source. For each view angle, two projection images are collected using filters whose absorption edge energies are different. The ratio of the two images yields the attenuation information for the elements whose absorption edges lie between the two filter energies, as indicated in **Figure 45**. A series of such image pairs can provide separate information on the spatial distribution of many elements. **Figure 46** shows an example in which the use of filters has selected two of the three metal inserts in the sample. The use of the filters reduces the already low intensity from the x-ray tube, and the use of the ratio of the two images presents a further limitation on the statistical quality of the projections. It is therefore important to use a small number of views to obtain the best possible projection images. In this example, 12 projections with full 3D rotation of axes were obtained. **Figure 47** shows the artefacts present in the reconstruction when the same number of views is obtained with single-axis rotation.

The electron microscope produces images in which contrast is due to attenuation, and a series of views at different angles can be reconstructed to show 3D structure. The use of an arbitrary series of angles is quite difficult to achieve for materials specimens because of diffraction of the electrons from planes of atoms in the crystal structure. This source of contrast is not easily modeled by the usual attenuation calculation since one voxel may have quite different values in different directions. For noncrystalline materials such as biological specimens, however, the reconstruction is straightforward (Engel & Massalski, 1984; Hegerl, 1989).

Figure 45. Diagram of the use of balanced absorption edge filters to isolate a single energy band. The plots show the absorption coefficient as a function of energy for two different filters containing the elements chromium and nickel. Elements in the sample with absorption edges between these two energies, such as manganese, iron, and cobalt, will be imaged in the ratio of the two intensities.






Figure 46. Reconstruction of the same three planes as shown in **Figure 44**, but using images obtained as a difference between two projections through different filters (Cr and Fe metal foils), which form a bandpass filter to select a narrow band of x-ray energies. Note that one of the inserts (Cr) is missing but the Mn and Fe inserts are visible. Reconstructed using 12 cone beam projections with rotations in 3D.




Figure 47. Reconstruction of the same plane as shown in **Figure 46**, reconstructed in the same way using two filters, but using 12 radial projections (rotating the sample about one axis only). Note the artefacts between and within the inserts.

Even more efficient than collecting a series of different views using multiple orientations of a single specimen is using images of many different but identical specimens that happen to have different orientations, as mentioned above. **Figure 48** shows an example. The two-dimensional image is an electron micrograph of a single virus particle. The specimen is an adenovirus that causes respiratory ailments.

The low dose of electrons required to prevent damage to the specimen makes the image very noisy. In a typical specimen, however, many such particles are present, each in a different, essentially random orientation. Collecting the various images, indexing the orientation of each image by referring to the location of the triangular facets on the virus surface, and performing a reconstruction produces a 3D reconstruction of the particle in which each voxel value is the electron density. Modeling the surface of the outer protein coat of the virus produces the surface-rendered image shown in **Figure 49** (Stewart and Burnett, 1991).

At a very different scale, tomography has also been performed on the earth itself using seismography. Seismic waves are created by earthquakes or large explosions such as nuclear weapons tests. Such large-magnitude events generate two types of waves that propagate through the earth to receivers (seismographs) at many different locations. P-waves (pressure waves) are compressional pulses which can penetrate through every part of the earth's interior, while S-waves (shear waves) are transverse deformations that cannot propagate through the liquid core. In fact, the presence of a liquid core was deduced in 1906 by the British seismologist, R. D. Oldham, from the shadow cast by the core in seismic S-wave patterns.

Figure 48. Transmission electron microscope image of single adenovirus particle.





Figure 49. Reconstruction of the adenovirus particle from many transmission images.

The paths of seismic waves are not straight (**Figure 50**), but bend because of the variations in temperature, pressure and composition within the earth which affect the speed of transmission just as the index of refraction of glass affects light and causes it to bend in a lens system. Also similar to the behavior of light, the seismic waves may reflect at interfaces where the speed of propagation varies abruptly. This happens at the core-mantle boundary and the surface of the inner core. The propagation velocities of the P- and S-waves are different, and respond differently to composition.

Figure 50. Diagram of paths taken by pressure and shear waves from earthquakes, which reveal information the density along the paths through the core and mantle, and the location of discontinuities.





Figure 51. Computed tomogram of the mantle, showing rock densities (light shades are hot, light rocks which are rising, and conversely).

Collecting many seismograms from different events creates a set of ray paths that do not uniformly cover the earth, but rather depend on the chance and highly nonuniform distribution of earthquakes and the distribution of seismographs. Nevertheless, analysis of the travel times of waves that have taken different paths through the earth permits forming a tomographic reconstruction. The density of the material (shown by shading in **Figure 51**) indicates the temperature and the direction of motion (cool, dense material is sinking through the mantle toward the core, while hot, light material is rising). Convection in the mantle is the driving force behind volcanism and continental drift.

Also of great utility are waves that have reflected (one or more times) from the various surfaces. For instance, the difference in travel times of S-waves that arrive directly vs. those which have reflected from the core-mantle boundary permit mapping the elevation of that boundary with a resolution better than 1 kilometer, and reveal that the boundary is not a smooth spherical surface. Because the relatively viscous mantle is floating on a low velocity liquid core, and it is the relatively fast motion of the latter that produces the earth's magnetic field, the study of this interface is important in understanding the earth's dynamics.

Global tomographic reconstruction is generally insensitive to the small details of structure such as faults, but another ongoing program to perform high-resolution tomography under the state of California (where many faults exist, which are of more than casual interest to the surface-dwelling humans) employs an array of high sensitivity seismographs, and uses the very frequent minor earthquakes there to map out the faults through the reflections that they produce.


High resolution tomography

Medical tomography has a typical resolution of about 1 mm, which is adequate for its purpose, and radiologists generally feel comfortable with a series of planar section images in standardized orientations in which they have been trained to recognize normal and abnormal features; but there is considerable interest in applying true 3D tomographic imaging to study the microstructure of various materials including metals, ceramics, composites and polymers, as well as larger industrial components. Some of the structural features cannot be determined from conventional 2D microscopy of cross section surfaces. This includes determining the number of particles of arbitrary or variable shape in a volume, and the topology of networks or pore structures, which control the permeability of materials to fluids.

This information can only be determined by having the full 3D data set, with adequate resolution, and ideally with cubic voxels. Resolution of the order of 1 μm has been demonstrated using a synchrotron as a very bright point source of x-rays. Resolution of about 10 μm is possible using more readily available sources such as microfocus x-ray tubes. Filtering such sources to produce element-specific imaging is also possible, as illustrated.

Cone-beam geometry is well suited to this type of microstructural imaging because it provides magnification of the structure (Johnson et al., 1986; Russ, 1988; Kinney et al., 1989, 1990; Deckman, 1989). **Figure 52** shows this schematically. The magnification is strictly geometric because x-rays

Figure 52. Diagram of a cone-beam imaging system.
The projection image magnification is the ratio of b/a . The attainable resolution is limited by the spot size of the x-ray source.






Figure 53. Twenty individual planes of reconstructed voxels showing a sintered alumina ceramic consisting of 100- μm diameter spheres.

are not refracted by lenses, but can amount to as much as 100:1. The projected images can be collected using conventional video technology, after conversion to visible light by a phosphor or channel plate and suitable intensification. The intensity of conventional small-spot, x-ray sources is very low, so the use of high-brightness sources, such as those available at a synchrotron, is

particularly desirable for high-resolution imaging, as is image averaging, which may be done by using the same cooled charge coupled device (CCD) cameras used for astronomical imaging.

As discussed in Chapter 10, 3D imaging requires many voxels, and the reconstruction process is very computer intensive. The time required to perform the reconstruction is, however, still shorter than that required to collect the various projection images. These images are generally photon-limited, with considerable noise affecting the reconstruction as indicated above. In order to collect reasonable-quality projections from a finite intensity source, the number of view angles is limited. Ideally the views should be arranged to cover the polar angles optimally in 3D space. This arrangement of course places demands on the quality of the mechanism used to perform the rotations and tilts, because the center of rotation must be constant and located within a few micrometers to preserve the image quality as discussed above. Helical scanning is usually easier to accomplish, and more commonly used.

The presentation of 3D information requires extensive use of computer graphics methods as shown in Chapter 10. **Figure 53** shows a simple series of planes of voxels from a 3D tomographic reconstruction of a porous alumina ceramic. The individual particles are approximately 100 μm diameter spheres that fill about 60% of the volume of the sample. The voxels are 10 μm cubes. **Figure 54** shows one of the projection sets through this specimen, a two-dimensional image in which the spherical particles overlap along the lines of sight and are partially transparent. A three-dimensional presentation of this data is shown in **Figure 55**.

Since the first edition of this book (1990) showed these examples from our own experimental setup, several commercial implementations of such instruments have become available (Wang and Vannier, 2001). Figure 56 shows images of bone that have approximately 10 μm resolution, reconstructing up to 1000 slices of 1000×1000 voxels in a 1 mm^3 specimen. Reconstructions that in 1990 required a Cray supercomputer are now being performed on a standard desktop machine.





Figure 54. A single 2D projection set through the structure shown in Figure 53.

Figure 55. Three-dimensional presentation of the data from Figure 53.



il-
or
es
is




Figure 56. Microtomography of bone:

(a) one projected view image;

(b–d) reconstructed transverse slices at the positions marked in image a (images courtesy Skyscan, Aartselaar, Belgium).

3D Image Visualization

Sources of 3D data

T rue three-dimensional (3D) imaging is becoming more accessible with the continued development of instrumentation. Just as the pixel is the unit of brightness measurement for a two-dimensional (2D) image, the voxel (volume element, the 3D analog of the pixel or picture element) is the unit for 3D imaging; and just as processing and analysis is much simpler if the pixels are square, so the use of cubic voxels is preferred for three dimensions, although it is not as often achieved.

Several basic approaches are used for volume imaging. In Chapter 11, 3D imaging by tomographic reconstruction was described. This is perhaps the premier method for measuring the density and in some cases the composition of solid specimens. It can produce a set of cubic voxels, although that is not the only or even the most common way that tomography is presently used. Most medical and industrial applications produce one or a series of 2D section planes, which are spaced farther apart than the lateral resolution within the plane (Baba et al., 1984, 1988; Briarty and Jenkins, 1984; Johnson and Capowski, 1985; Kriete, 1992).

Tomography can be performed using a variety of different signals, including seismic waves, ultrasound, magnetic resonance, conventional x-rays, gamma rays, neutron beams, and electron microscopy, as well as other even less familiar methods. The resolution may vary from kilometers (seismic tomography), to centimeters (most conventional medical scans), millimeters (typical industrial applications), micrometers (microfocus x-ray or synchrotron sources), and even nanometers (electron microscope reconstructions of viruses and atomic lattices). The same basic presentation tools are available regardless of the imaging modality or the dimensional scale.

The most important variable in tomographic imaging, as for all of the other 3D methods discussed here, is whether the data set is planes of pixels, or an array of true voxels. As discussed in Chapter 11, it is possible to set up an array of cubic voxels, collect projection data from a series of views in three dimensions, and solve (either algebraically or by backprojection) for the density of each voxel. The most common way to perform tomography however, is to define one plane at a time as an array of square pixels, collect a series of linear views, solve for the 2D array of densities in that plane, and then proceed to the next plane. When used in this way, tomography shares

many similarities (and problems) with other essentially 2D imaging methods that we will collectively define as serial imaging or serial section techniques.

A radiologist viewing an array of such images is expected to combine them in his or her mind to "see" the 3D structures present. (This process is aided enormously by the fact that the radiologist already knows what the structure is, and is generally looking for things that differ from the familiar, particularly in a few characteristic ways that identify disease or injury.) Only a few current-generation systems use the techniques discussed in this chapter to present 3D views directly. In industrial tomography, the greater diversity of structure (and correspondingly lesser ability to predict what is expected), and the greater amount of time available for study and interpretation, has encouraged the use of computer graphics. But such displays are still the exception rather than the rule, and an array of 2D planar images is more commonly used for volume imaging. This chapter emphasizes methods that use a series of parallel, uniformly spaced 2D images, but present them in combination to show 3D structure.

These images are obtained by dissecting the sample into a series of planar sections, which are then piled up as a stack of voxels. Sometimes the sectioning is physical. Blocks of embedded biological materials, textiles, and even some metals, can be sliced with a microtome, and each slice imaged (just as individual slices are normally viewed). Collecting and aligning the images produces a 3D data set in which the voxels are typically very elongated in the "Z" direction because the slices are much thicker or more widely spaced than the lateral resolution within each slice.

At the other extreme, the secondary ion mass spectrometer uses an incident ion beam to remove one layer of atoms at a time from the sample surface. These pass through a mass spectrometer to select atoms from a single element, which is then imaged on a fluorescent screen. Collecting a series of images from many elements can produce a complete 3D map of the sample. One difference from the imaging of slices is that there is no alignment problem, because the sample block is held in place as the surface layers are removed. On the other hand, the erosion rate through different structures can vary so that the surface does not remain planar, and this roughening or differential erosion is very difficult to account for. In this type of instrument, the voxel height can be very small (essentially atomic dimensions) while the lateral dimension is many times larger.

Serial sections

Most physical sectioning approaches are similar to one or the other of these examples. They are known collectively as serial section methods. The name serial section comes from the use of light microscopy imaging of biological tissue, in which blocks of tissue embedded in resin are cut using a microtome into a series of individual slices. Collecting these slices (or at least some of them) for viewing in the microscope enables researchers to assemble a set of photographs which can then be used to reconstruct the 3D structure.

This technique illustrates most of the problems that may be encountered with any 3D imaging method based on a series of individual slices. First, the individual images must be aligned. The microtomed slices are collected on slides and viewed in arbitrary orientations. So, even if the same structures can be located in the different sections (not always an easy task, given that some variation in structure with depth must be present or there would be no incentive to do this kind of work), the pictures do not line up.

Using the details of structure visible in each section provides only a coarse guide to alignment. The automatic methods generally seek to minimize the mismatch between sections either by aligning the centroids of features in the planes so that the sum of squares of distances is minimized, or by overlaying binary images from the two sections and shifting or rotating to minimize the area resulting from combining them with an Ex-OR (exclusive OR) operation, discussed in Chapter 7.

This procedure is illustrated in **Figure 1**. When grey-scale values are present in the image, cross-correlation can be used as discussed in Chapter 5. Unfortunately, neither of these methods is easy to implement in the general case when sections may be shifted in X and Y and also rotated. Solving for the “best alignment” is difficult and must usually proceed iteratively and slowly.

Furthermore, there is no reason to expect the minimum point reached by these algorithms to really represent the true alignment. As shown in **Figures 2** and **3**, shifting or rotating each image to visually align the structures in one section with the next can completely alter the reconstructed 3D structure. It is generally assumed that given enough detail present in the images, some kind of average alignment will avoid these major errors; however, it is far from certain that a best visual alignment is the correct one nor that automated methods, which overlap sequential images, produce the proper alignment.

One approach that improves on the use of internal image detail for alignment is to incorporate fiducial marks in the block before sectioning. These could take the form of holes drilled by a laser, threads or fibers placed in the resin before it hardens, or grooves machined down the edges of the block, for example. With several such marks that can reasonably be expected to maintain their shape from section to section and continue in some known direction through the stack of images, much better alignment is possible. Placing and finding fiducial marks in the close vicinity of the structures of interest is often difficult. In practice, if the sections are not contiguous there may still be difficulties, and alignment errors may propagate through the stack of images.

Most fiducial marks are large enough to cover several pixels in each image. As discussed in Chapter 8, this size allows locating the centroid to a fraction of one pixel accuracy, although not all systems take advantage of this capability. Once the alignment points are identified (either from fiducial marks or internal image detail), the rotation and translation of one image to line up with the next is performed as discussed in Chapter 3. Resampling of the pixel array and interpolation to prevent aliasing produces a new image. This process takes some computational time, but this is a minor problem in comparison to the difficulty of obtaining the images in the first place.

Unfortunately, for classic serial sectioning the result of this rotation and translation is not a true representation of the original 3D structure. The act of sectioning using a microtome generally produces

Figure 1. Alignment of serial sections by “best fit” of features seeks to minimize mismatched area, measured by ExOR function, as a function of translation and rotation.




Figure 2. Alignment of serial sections with translation.
Sections through an inclined circular cylinder may be misconstrued as a vertical elliptical cylinder.





Figure 3. Alignment of serial sections with rotation:
(a) actual outlines in 3D serial section stack;
(b) surface modeling applied to outlines, showing twisted structure;
(c) erroneous result without twist when outlines are aligned to each other.

some distortion in the block. This 5–20% compression in one direction is usually assumed to be nearly the same for all sections (since they are cut in the same direction, and generally have only small differences in structure that would alter their mechanical properties). If the fiducial marks have known absolute coordinates, then stretching of the images to correct for the distortion is possible. It is usually assumed that the entire section is compressed uniformly, although for some samples this may not be true.

Otherwise, it may be possible to use internal information to estimate the distortion. For example, if there is no reason to expect cells or cell nuclei to be elongated in any preferred direction in the tissue, then measurement of the dimensions of many cells or nuclei may be used to determine an average amount of compression. Obviously, this approach includes some assumptions and can only be used in particular circumstances.

Another difficulty with serial sections is calibration of dimension in the depth direction. The thickness of the individual sections is only known approximately (for example, by judging the color of the light produced by interference from the top and bottom surfaces, or based on the mechanical feed rate of the microtome). It may vary from section to section, and even from place to place within the section, depending on the local hardness of the material being cut. Constructing an accurate depth scale is quite difficult, and dimensions in the depth direction will be much less accurate than those measured within one section plane.

If only some sections are used, such as every second or fifth (in order to reduce the amount of work required to image them and then align the images), then this error becomes much worse.

It also becomes difficult to follow structures from one image to the next with confidence. Before computer reconstruction methods became common, however, this kind of skipping was often necessary simply to reduce the amount of data that the human observer had to juggle and interpret.

Using only a fraction of the sections is particularly common when ultra-thin sections are cut for viewing in an electron microscope instead of the light microscope. As the sections become thinner, they increase in number and are more prone to distortion. Some may be lost (for instance due to folding) or intentionally skipped. Portions of each section are obscured by the support grid, which also prevents some from being used. At higher magnification, the fiducial marks become larger, less precisely defined, and above all more widely spaced so that they may not be in close proximity to the structure of interest.

Figure 4 shows a portion of a series of transmission electron microscope (TEM) images of tissue in which the 3D configuration of the membranes (dark, stained lines) is of interest. The details of the edges of cells and organelles have been used to approximately align pairs of sections through the stack, but different details must be used for different pairs as there is no continuity of detail through the entire stack. The membranes can be isolated in these images by thresholding (**Figure 5**), but the sections are too far apart to link the lines together to reconstruct the 3D shape of the surface. This problem is common with conventional serial section images.

Metallographic imaging typically uses reflected rather than transmitted light. As discussed in the next section, serial sectioning in this context is accomplished by removing layers of materials sequentially by physical polishing. The need to locate the same sample position after polishing, and to monitor the depth of polishing, can be met by placing hardness indentations on the sample, or by laser ablation of pits. These serve as fiduciary marks for alignment, and the change in size of the mark reveals the depth. In archaeological excavation the fiduciary marks may be a network of strings and a transit, and the removal tool may be a shovel. In some mining and quarrying examples it may be a bulldozer, but the principles remain the same regardless of scale.





Figure 4. Four serial section images from a stack (courtesy Dr. C. D. Bucana, University of Texas M. D. Anderson Cancer Center, Houston, TX), which have already been rotated for alignment. The membranes at the upper left corner of the images are thresholded and displayed for the entire stack of images in **Figure 5**.

Figure 5. Membranes from the sequential images illustrated in **Figure 4**, showing the changes from section to section. Because of the separation distance between the planes, these variations are too great to model the shape of the surfaces in 3D.



In most of these surface imaging methods, the sequential 2D images represent the sample along planes that are separated in the z direction. The intervening material that has been removed must be inferred or interpolated from the planes. In the cases, the voxel value is not really an average over its extent, as most pixel values are. Nor is the voxel truly a discrete point value in space, because it does represent an average in the section plane. Interpolation between sections that are too far apart (in terms of the scale of the structure) can lead to some serious errors and misinterpretation.

For the serial sectioning method in which slices are viewed in transmission the voxel value is a volume average, which is easier to interpret. In most cases, the voxel value is a measure of density of the material. Depending on what the radiation is (visible light, x-rays, electrons, neutrons, sound, and so forth), the value may represent the local concentration of some element or compound. In some cases, emitted radiation from voxels also gives concentration information (examples include fluorescence light microscopy and positron-emission tomography).

Optical sectioning

Physical sectioning on any scale is a difficult technique that destroys the sample. Controlling and measuring the section thickness and aligning the sections or at least locating the same position on the sample can become a major source of error. In some cases, it is possible to image sections through a sample without performing physical sectioning. The confocal scanning light microscope (CSLM) offers one way to accomplish this. Depth information can also be obtained in some cases by interpreting the phase shift introduced by objects viewed in light microscopy, as well as by neutron or x-ray transmission imaging (Barty et al., 2000).

The normal operation of the transmission light microscope does not lend itself to optical sectioning. The depth of field of high numerical aperture optics is small (just a few times the lateral resolution), so that only a small "slice" of the image will be sharply focused; however, light from locations above and below the plane of focus is also transmitted to the image, out of focus, and this both blurs the image and includes information from an extended distance in the z direction. In some cases, deconvolution of the point spread function can be accomplished (as discussed in Chapter 5), but in many cases this is of limited value since the blurring varies from one location to another. An example shown in the section on "stereo viewing," (**Figure 36**) under the topic of stereo imaging shows that it is possible to process the images to remove some of the artefacts that result from the passage of light through the sample above and below the plane of focus.

The confocal microscope eliminates this extraneous light, and so produces useful optical section images without the need for processing. This is possible because the sample is imaged one point at a time (thus, the presence of “scanning” in the name). The principle was introduced in Chapter 4 (Image Enhancement), in conjunction with some of the ways that images of light reflected from surfaces can be processed. The principle of the confocal microscope is that light from a point source (often a laser) is focused on a single point in the specimen and collected by an identical set of optics, reaching a pinhole detector. Any portion of the specimen away from the focal point, and particularly out of the focal plane, cannot return light to the pinhole light to interfere with the formation of the image. Scanning the beam with respect to the specimen (by moving the light source, the specimen, or using scanning elements in the optical path) builds up a complete image of the focal plane.

If the numerical aperture of the lenses is high, the depth of field of this microscope is very small, although still several times the lateral resolution within individual image planes. Much more important, the portion of the specimen that is away from the focal plane contributes very little to the image. This makes it possible to image a plane within a bulk specimen, even one that would ordinarily be considered translucent because of light scattering. This method of isolating a single plane within a bulk sample, called optical sectioning, works because the confocal light microscope has a very shallow depth of field and a high rejection of stray light. Translating the specimen in the z direction and collecting a series of images makes it possible to build up a three-dimensional data set for viewing.

Several imaging modalities are possible with the confocal light microscope. The most common are reflected light, in which the light reflected from the sample returns through the same objective lens as used to focus the incident light and is then diverted by a mirror to a detector, and fluorescence in which light is emitted from points within the specimen and is recorded using the same geometry. It is also possible, however, to use the microscope to view transmitted light images using the geometry shown in **Figure 6**. This permits acquiring transmitted light images for focal plane sectioning of bulk translucent or transparent materials. **Figure 7** shows an example of a transmitted light focal plane section.

Both transmitted and reflected-light images of focal plane sections can be used in 3D imaging for different types of specimens. The characteristic of reflected-light confocal images is that the intensity of light reflected to the detector drops off very rapidly as points are shifted above or below the focal plane. Therefore, for structures in a transparent medium, only the surfaces will reflect light.




Figure 6. Transmission confocal scanning light microscopy (CSLM) can be performed by passing the light through the specimen twice. Light is not imaged from points away from the in-focus point, which gives good lateral and excellent depth resolution compared to a conventional light microscope. (The more common reflected light confocal microscope omits the optics beneath the specimen.)

Figure 7. CSLM image showing a $\frac{1}{30}$ second image of a paramecium swimming in a droplet of water, as it passed through the focal plane of the microscope.



For any single image plane, only the portion of the field of view where some structure passes through the plane will appear bright, and the rest of the image will be dark. This characteristic permits some rather straightforward reconstruction algorithms.

A widely used imaging method for the confocal microscope is emission or fluorescence, in which the wavelength of the incident light is able to cause excitation of a dye or other fluorescing probe introduced to the specimen. The lower-energy (longer wavelength) light emitted by this probe is separated from the incident light, for instance by a dichroic mirror, and used to form an image in which the location of the probe or dye appears bright. Building up a series of images in depth allows the structure labeled by the probe to be reconstructed.

The transmitted-light mode, while it is the most straightforward in terms of optical sectioning, is little used as yet. This situation is partly due to the difficulties in constructing the microscope with matched optics above and below the specimen, as compared with the reflection and emission modes in which the optics are only above it; however, the use of a lens and mirror beneath the specimen to return the light to the same detector as present in the more standard microscope design can produce most of the same imaging advantages (the only loss is that in passing through the specimen twice, some intensity is lost).

The principal advantages of optical sectioning are avoiding physical distortion of the specimen due to cutting, and having alignment of images from the various imaging planes. The depth resolution, although inferior to the lateral resolution in each plane by about a factor of two to three times, is still useful for many applications. This difference in resolution, however, does raise some difficulties for 3D image processing, even if the distance between planes is made smaller than the resolution so that the stored voxels are cubic (which is by no means common).

By measuring or modeling the 3D shape of the microscope's point spread function, it is possible by deconvolution to improve the resolution of the confocal light microscope. The method is identical to that shown in Chapter 5 for 2D images.

Sequential removal

Many materials are opaque and therefore cannot be imaged by any transmission method, preventing any type of optical sectioning. Indeed, metals, composites, and ceramics are usually examined in the reflected light microscope, although it is still possible to collect a series of depth images for 3D reconstruction by sequential polishing of such materials, as mentioned earlier.

The means of removal of material from the surface depends strongly on the hardness of the material. For some soft metals, polymers, and textiles, the microtome can be used just as for a block of biological material, except that instead of examining the slice of material removed, the surface left behind is imaged. This approach avoids most problems of alignment and distortion, especially if the cutting can be done *in situ* without removing the specimen from the viewing position.

in the microscope. It is still difficult to determine precisely the thickness of material removed in each cut and to assure its uniformity, but it is generally estimated from the mechanical settings on the device (ignoring any permanent or temporary distortion in the material).

For harder materials, the grinding or polishing operations used to produce conventional sections for 2D images can be used. Such operations generally require removing and replacing the specimen, so again fiducial marks are needed to locate the same region. Probably the most common approach to this marking is the use of hardness indentations. Several pyramid-shaped impressions are made in the surface of the specimen so that after additional abrasion or polishing, the deepest parts of the indentations are still visible. These can be accurately aligned with the marks in the original image. In addition, the reduction in size of the impression, with a shape that is known, gives a measure of the depth of polish and hence of the spacing between the two images. With several such indentations, the overall uniformity of polish can also be judged, although local variations due to the hardness of particular phases may be present.

For still harder materials or ones in which conventional polishing might cause surface damage, other methods may be used. Electrolytic or chemical etching is generally difficult to control and little used. Ion beam erosion is slow, but is already in use in many laboratories for the cutting and thinning of transmission electron microscope specimens, and may be utilized for this purpose. Controlling the erosion to obtain uniformity and avoid surface roughening presents challenges for many specimens.

In situ ion beam erosion is used in the scanning electron microscope and scanning Auger microscope, for instance to allow the removal of surface contamination. Focused ion beams (FIB) are also used to cut slices in the *z* direction to examine microstructures. This capability can be used to produce a series of images in depth in these microscopes, which generally have resolution far better than the light microscope. The time involved in performing the erosion or slicing may be quite long (and thus costly), however, and the uniformity of eroding through complex structures (the most interesting kind for imaging) may be poor.

One kind of microscope performs this type of erosion automatically as part of its imaging process. The ion microscope or secondary ion mass spectrometer (SIMS) uses a beam of heavy ions to erode a layer of atoms from the specimen surface. The secondary ions are then separated according to element in a mass spectrometer and recorded, for example using a channel plate multiplier and more-or-less conventional video or digital camera, to form an image of one plane in the specimen for one element at a time. The depth of erosion is usually calibrated for these instruments by measuring the signal profile of a known standard, such as may be produced by the same methods used to produce modern microelectronics.

The rate of surface removal is highly controllable (if somewhat slow) and capable of essentially atomic resolution in depth. The lateral resolution, by contrast, is of about the same level as in the conventional light microscope, so in this case instead of having voxels which are high in resolution in the plane but poorer in the depth direction, the situation is reversed. The non-cubic voxels create problems for processing and measurement.

Furthermore, the erosion rate for ion beam bombardment in the ion microscope or SIMS may vary from place to place in the specimen as a function of composition, structure, or even crystallographic orientation. This variation does not necessarily show up in the reconstruction, since each set of data is assumed to represent a plane, but can cause significant distortion in the final interpretation. In principle, stretching of the data in 3D can be performed just as images can be corrected for deformation in 2D, although without fiducial marks or accurate quantitative data on local erosion rates, it is hard to accomplish this with real data.

The ability to image many different elements with the SIMS creates a rich data set for 3D display. A color 2D image has three channels (whether it is saved as RGB or HSI, as discussed in Chapter 1), and satellite 2D images typically have as many as seven bands including infrared. The SIMS may have practically any number. The ability of the instrument to detect trace levels (typically ppm or better) of every element or even isotope in the periodic table, plus molecular fragments, means that even for relatively simple specimens the multiband data present a challenge to store, display and interpret.

Another type of microscope that removes layers of atoms as it images them is the atom probe ion microscope. In this instrument, a strong electrical field between a sharply curved sample tip and a display screen causes atoms to be desorbed from the surface and accelerated toward the screen where they are imaged. The screen may include an electron channel plate to amplify the signal so that individual atoms can be seen, or may be used as a time-of-flight mass spectrometer with pulsed application of the high voltage so that the different atom species can be distinguished. With any of the instrument variations, the result is a highly magnified image of atoms from the sample, showing atom arrangements in 3D as layer after layer is removed.

Examples of images from all these types of instruments were shown in Chapter 1.

Stereo measurement

There remains another way to see 3D structures. It is the same way that humans see depth in some real-world situations — having two eyes that face forward so that their fields-of-view overlap permits us to use stereoscopic vision to judge the relative distance to objects. In humans, this is done point by point, by moving our eyes in their sockets to bring each subject to the fovea, the portion of the retina with the densest packing of cones. The muscles in turn tell the brain what motion was needed to achieve convergence, and so we know whether one object is closer or farther than another.

Further into this section, we will see stereo vision used as a means to transmit 3D data to the human viewer. It would be wrong to think that all human depth perception relies on stereoscopy. In fact, much of our judgment about the 3D world around us comes from other cues such as shading, relative size, precedence, atmospheric effects (e.g., fog or haze), and motion flow (nearer objects move more in our visual field when we move our head) that work just fine with one eye and are used in some computer-based measurement methods (Roberts, 1965; Horn, 1970, 1975; Woodham, 1978; Carlsen, 1985; Pentland, 1986). For the moment, however, let us see how stereoscopy can be used to determine depth information to put information into a 3D computer database.

The light microscope has a rather shallow depth of field, which is made even less in the confocal scanning light microscope discussed previously. Consequently, looking at a specimen with deep relief is not very satisfactory except at relatively low magnifications; however, the electron microscope has lenses with very small aperture angles, and thus has very great depth of field. Stereoscopy is most commonly used with the scanning electron microscope (SEM) to produce in-focus images of rough surfaces. Tilting the specimen, or electromagnetically deflecting the scanning beam, can produce a pair of images from different points of view that form a stereo pair. Looking at one picture with each eye fools the brain into seeing the original rough surface.

Measuring the relief of surfaces from such images is the same in principle and in practice as using stereo pair images taken from aircraft or satellites to measure the elevation of topographic features on the earth or another planet. The richer detail in the satellite photos makes it easier to find matching points practically anywhere in the images, but by the same token requires more matching points to define the surface than the simpler geometry of typical specimens observed in the SEM. The mathematical relationship between the measured parallax (the apparent displacement of points in the left and right eye image) and the relative elevation of the two points on the surface was presented in Chapter 1.

Automatic matching of points from stereo pairs is a difficult task for computer-based image analysis (Marr and Poggio, 1976; Medioni and Nevatia, 1985; Kayaalp and Jain, 1987). It is usually performed by using the pattern of brightness values in one image, for instance the left one, as a template to perform a cross-correlation search for the most nearly identical pattern in the right image. The area of search is restricted by the possible displacement, which depends on the angle between the two views and the maximum roughness of the surface, to a horizontal band in the second image. Some points will not be matched by this process because they may not be visible in both images (or are lost off the edges of one or the other image). Other points will match poorly because the local pattern of brightness values in the pixels includes some noise, and several parts of the image may have similar noise levels.

Matching many points produces a new image in which each pixel can be given a value based on the parallax, and hence represents the elevation of the surface. This range image will contain many false matches, but operations such as a median filter usually do a good job of removing the outlier points to produce an overall range image of the surface. In the example of **Figure 8**, cross-correlation matching of every point in the left-eye view with points in the right produces a disparity map (the horizontal distance between the location of the matched points) that contains false matches, which are filled in by a median filter as shown. The resulting elevation data can be used for measurement of points or line profiles, or used to reconstruct surface images, as illustrated. This use of surface range data is discussed further in Chapter 13.



a



b

Figure 8. Elevation measurement using stereo pair images:

(**a, b**) left and right eye views of a microfossil;

(**c**) raw cross-correlation disparity values;

(**d**) median filter applied to **c**;

(**e**) surface height values measured from a mean plane, displayed as grey-scale values;

(**f**) rendered perspective-corrected surface model using the values in **e**;


(**g**) the same surface reconstruction using elevation values from **e** with surface brightness values from **a**.




c



d



e



f



g

A second approach to matching stereo pairs is based on the realization that many of the points in each image will not match well because they are just part of the overall surface or structure and not the “interesting” points where surfaces meet or other discontinuities are present. This approach is presumably related to human vision, which usually spends most of its time concentrating on only the few points in each scene where discontinuities are found. Locating these interesting points based on some local property such as the variance, entropy, or result of a high-pass filter, produces a comparatively short list of points to be matched between the two images (Moravec, 1977; Quam and Hannah, 1974). A typical case may have only thousands of points, instead of the million or so pixels in the original images.

Somewhat better results are obtained by fitting a cubic polynomial to the pixel brightness values in each neighborhood. Then, if the polynomial is written as

$$f(i, j) = c_1 + c_2x + c_3y + c_4x^2 + c_5xy + c_6y^2 + c_7x^3 + c_8x^2y + c_9xy^2 + c_{10}y^3 \quad (1)$$

the Zuniga–Haralick operator (Zuniga and Haralick, 1983; Haralick and Shapiro, 1992) used to detect corner points is

$$\frac{-2 \cdot (c_2^2c_6 - c_2c_3c_5 - c_3^2c_4)}{(c_2^2 + c_3^2)^{3/2}} \quad (2)$$

The points on the resulting short list are then matched in the same way as above, by correlation of their neighborhood brightness patterns. Additionally, for most surfaces the order of points is preserved. This, and the limits on possible parallax for a given pair of images, reduces the typical candidate list for matching to ten or less, and produces a list of surface points and their elevations. It is then assumed that the surface between these points is well behaved and can be treated as consisting of planar facets or simple spline patches, which are constructed by linking the points in a Delaunay tessellation. If the facets are small enough, it is possible to generate a contour map of the surface as shown in **Figure 9** by interpolating straight line segments between points along the edges of each planar facet. A complete display of elevation, called a range image, can be produced by interpolation as shown in **Figure 10**.

The TEM also has a very large depth of field. In most cases, the specimens observed in the TEM are very thin (in order to permit electron penetration), and the optical depth of field is unimportant. With recent generations of high-voltage microscopes, however, comparatively thick samples (of the order of micrometers) may be imaged. This thickness is enough to contain a considerable amount of 3D structure at the resolution of the TEM (of the order of a few nanometers). Thus, using the same approach of tilting the specimen to acquire stereo pair images, it is possible to obtain information about the depth of points and the 3D structure.

Figure 9. Drawing contour lines (iso elevation lines) on the triangular facets joining an arbitrary arrangement of points with elevations that have been determined by stereoscopy.






Figure 10. Interpolation of a range image:
(a) isolated, randomly arranged points with measured elevation (color coded);
(b) contour lines drawn through the tessellation;
(c) smooth interpolation between contour lines;
(d) the constructed range image (grey scale).


Presenting the images to a human viewer's eyes so that two pictures acquired at different times can be fused in the mind and examined in depth is not difficult. It has been accomplished for years photographically, and is now often done with modest tradeoff in lateral resolution using a computer to record and display the images. The methods discussed in the following paragraphs, which use stereo pair displays to communicate 3D information from generated images are equally applicable here.

It is far more difficult, however, to have the computer determine the depth of features in the structure and construct a 3D database of points and their relationship to each other. Part of the problem is that so much background detail is available from the (mostly) transparent medium surrounding the features of interest that it may dominate the local pixel brightnesses and make matching impossible. Another part of the problem is that it is no longer possible to assume that points maintain their order from left to right. In a 3D structure, points may change their order as they pass in front or in back of each other.

The consequence of these limitations has been that only in a very few, highly idealized cases has automatic fusion of stereo pair images from the TEM been attempted successfully. Simplification of the problem using very high contrast markers, such as small gold particles bound to selected surfaces using antibodies, or some other highly selective stain, helps. In this case, only the markers are considered. Only a few dozen of these exist, and similar to the interesting points mentioned previously for mapping surfaces, they are easily detected (being usually far darker than anything else in the image) and only a few could possibly match.

Even with these markers, a human may still be needed to identify the matches. Given the matching points in the two images, the computer can construct a series of lines that describe the surface which the markers define, but this surface may be only a small part of the total structure. **Figure 11** shows an example of this method in which human matching was performed. Similar methods can be applied to stained networks (Huang et al., 1994), or the distribution of precipitate particles in materials, for example.

Figure 11. Example of decorating a surface with metal particles (Golgi stain) shown in transmission electron micrographs (Peachey and Heath; 1989) with elevations that are measured stereoscopically to form a network describing the surface (b).



In most matching procedures, the points in left and right images are defined in terms of pixel addresses. The error in the vertical dimension determined by stereoscopy is typically an order-of-magnitude greater than the precision of measurement of the parallax, because the vertical height is proportional to the lateral parallax times the cosecant of the small angle between the views. Improving the measurement of parallax between features to subpixel accuracy is therefore of considerable interest. Such improvement is possible in some cases, particularly when information from many pixels can be combined. As described in Chapter 9, the centroids of features or the location of lines can be specified to an accuracy of one-tenth of a pixel or better.

3D data sets

In the case of matching of points between two stereo pair images, the database is a list of a few hundred or perhaps thousands of coordinates that usually define either a surface or perhaps nodes in a network structure. If these points are to be used for measurement, the coordinates and perhaps some information on which points are connected to which points are all that is required. If image reconstruction is intended, it will be necessary to interpolate additional points between them to complete a display. This is somewhat parallel to the use of boundary representation in two dimensions. It may offer a very compact record of the essential (or at least of the selected) information in the image, but it requires expansion to be visually useful to the human observer.

The most common way to store 3D data sets is as a series of 2D images. Each single image, which we have previously described as an array of pixels, is now seen to have depth. This depth is present either because the plane is truly an average over some depth of the sample (as in looking through a thin section in the light microscope) or based on the spacing between that plane and the next (as for instance a series of polished planes observed by reflected light). Because of the depth associated with the planes, we refer to the individual elements as voxels (volume elements) rather than pixels (pixel elements).

For viewing, processing and measurement, the voxels will ideally be regular and uniformly spaced. This goal is often accomplished with a cubic array of voxels, which is easiest to address in computer memory and corresponds to the way that some image acquisition devices function (e.g., the confocal scanning light microscope). Other arrangements of voxels in space offer some advantages. In a simple cubic arrangement, the neighboring voxels are at different distances from the central voxel, depending on whether they share a face, edge or corner. Deciding whether voxels touch and are part of the same feature requires a decision to include 6-, 18- or 26-neighbor connectedness, even more complicated than the 4- or 8-connectedness of square pixels in a 2D image, discussed in Chapter 6.

More symmetrical arrangements are possible. The arrangements of atoms in metal crystals typically occupy sites in one of three lattice configurations: body-centered cubic (BCC), face-centered cubic (FCC), or hexagonal close packed (HCP). The first of these surrounds each atom (or voxel) with eight touching neighbors at the same distance, and the other two have 12 equidistant neighbors.

The advantage of these voxel-stacking arrangements is that processing of images can treat each of the neighbors identically, and that measurements are less biased as a function of direction. A more symmetrical neighborhood with neighbors at uniform distances also simplifies processing, including the application of filters and morphological operations. Of course, in order to fill space, the shapes of the voxels in these cases are not simple. Storing and addressing the voxel array is difficult, as is acquiring images or displaying them. Usually the acquired image must be re-sampled by interpolation to obtain voxels in one of these patterns, and a reverse interpolation is needed for display. For most purposes, these disadvantages outweigh the theoretical advantages, just as the use of a hexagonal pixel array is rarely used for two-dimensional images. Cubic arrays are the most common 3D arrangement of voxels.

If the voxels are not cubic because the spacing between planes is different from the resolution within each plane, it may be possible to adjust things so that they are. In the discussion that follows, we will assume that the depth spacing is greater than the spacing within the plane, but an analogous situation could be described for the reverse case. This adjustment could be done by interpolating additional planes of voxels between those that have been measured. Unfortunately, doing this will not help much with image processing operations, since the assumption is that all of the neighbors are equal in importance, and with interpolation they become redundant.

The alternative approach is to reduce the resolution within the plane by sampling every n th pixel, or perhaps by averaging pixels together in blocks, so that a new image is formed with cubic voxels. This resolution reduction also reduces the amount of storage that will be required, since many fewer voxels remain. Although it seems unnatural to give up resolution, in a few cases where cubic voxels are required for analysis this is done.

A variety of different formats are available for storing 3D data sets, either as a stack of individual images or effectively as an array of voxel values with x, y, z indices. Such arrays become very large, very fast. A 512×512 or 640×480 pixel 2D image, which represents a very common image size for digitizing video images, occupies 250 or 300K bytes of storage, using one byte per pixel (256 grey values). This is easily handled in the memory of a desktop computer, or recorded on a floppy disk. Even with modern digital cameras, a typical 1024×1024 color image occupies only a few M bytes. A $512 \times 512 \times 512$ 3D image would occupy 128 megabytes of memory, about the upper limit for practical processing on a desktop machine. Larger files present difficulties just in storing or transmitting from place to place, let alone processing. Most of the data sets in this chapter are much smaller, for instance $256 \times 256 \times 50$ (3.2 megabytes). The same operations shown here can be used for larger data sets, given time, computer power (speed and memory), or both.

Compression can reduce the size of the storage requirements. The individual 2D images can be compressed either using run-length encoding (especially useful for binary images), or by the JPEG (Joint Photographers Expert Group) compression algorithms based on a discrete cosine transform, discussed in Chapter 2. There is not yet a standardized algorithm extending to 3D images, but there are emerging standards for time sequences of images. One, the MPEG (Moving Pictures Expert Group) approach, is based on the fact that in a time series of images, most of the pixels do not change much in successive frames. Similar assumptions have been used to transmit video conferences over telephone lines; only the changing pixels must be transmitted for each frame.

This high correlation from frame to frame is also usually true for a series of planes in 3D, and may lead to standardized algorithms for compacting such images. They will still require unpacking for display, processing and measurement, however, and the significant artefacts introduced by these

lossy compression methods (discussed in Chapter 2) argues against their use. For 3D image stacks compressed using MPEG methods, the artefacts in the z -direction (interpreted as time) are even worse than those in each individual x - y plane.

It is instructive to compare this situation to that of computer-aided design (CAD). For manmade objects with comparatively simple geometric surfaces, only a tiny number of point coordinates are required to define the entire 3D structure. This kind of boundary representation is very compact, but it often takes some time (or specialized display hardware) to render a drawing with realistic surfaces from such a data set. For a voxel image, the storage requirements are great but information is immediately available without computation for each location, and the various display images shown in this chapter can usually be produced very quickly (sometimes even at interactive speeds) by modest computers.

For instance, given a series of surfaces defined by boundary representation or a few coordinates, the generation of a display may proceed by first constructing all of the points for one plane, calculating the local angles of the plane with respect to the viewer and light source, using those to determine a brightness value, and plotting that value on the screen. At the same time, another image memory is used to store the actual depth (z -value) of the surface at that point. After one plane is complete, the next one is similarly drawn except that the depth value is compared point by point to the values in the z -buffer to determine whether the plane is in front of or behind the previous values. Of course, each point is only drawn if it lies in front. This procedure permits multiple intersecting planes to be drawn on the screen correctly. (For more information on graphic presentation of 3D CAD data, see Foley and Van Dam, 1984; or Hearn and Baker, 1986.)

Additional logic is needed to clip the edges of the planes to the stored boundaries, to change the reflectivity rules used to calculate brightness depending on the surface characteristics, and so forth. Standard texts on computer graphics describe algorithms for accomplishing these tasks and devote considerable space to the relative efficiency of various methods because the time involved can be significant. By comparison, looking up the value in a large array, or even running through a column in the array to add densities or find the maximum value, is very fast. This is particularly true if the array can be held in memory rather than requiring disk access.

The difficulties of aligning sequential slices to produce a 3D data set were discussed above. In many cases, there may be several 3D data sets obtained by different imaging techniques (e.g., magnetic resonance imaging (MRI), x-ray, and PET images of the head) which must be aligned to each other. They also commonly have different resolutions and voxel sizes, so that interpolation is needed to adjust them to match one another. The situation is a direct analog to the 2D problems encountered in geographical information systems (GIS) in which surface maps, images in different wavelengths from different satellites, aerial photographs and other information must be aligned and combined.

The general problem is usually described as one of registering the multiple data sets. The two principal techniques, which are complementary, are to use cross-correlation methods on the entire pixel or voxel array as discussed in Chapter 5, or to isolate specific features in the multiple images and use them as fiducial marks to perform warping (Brown, 1992; Besl, 1992; van den Elsen et al., 1993, 1994, 1995; Reddy and Chatterji, 1996; Frederik et al., 1997; West et al., 1997).

Slicing the data set


Most 3D image data sets are actually stored as a series of 2D images, so it is very easy to access any of the individual image planes, usually called slices. Playing the series of slices back in order to create an animation or "movie" is perhaps the most common tool available to let the user view the data. It is often quite effective in letting the viewer perform the 3D integration, and as it

recapitulates the way the images may have been acquired (but with a much compressed time base), most viewers can understand images presented in this way. A simple user interface need only allow the viewer to vary the speed of the animation, change direction or stop at a chosen slice, for example. The same software now widely available to play back movies on the computer screen can be used for this purpose.

One problem with presenting the original images as slices of the data is that the orientation of some features in the 3D structure may not show up very well in the slices. It is useful to be able to change the orientation of the slices to look at any plane through the data, either in still or animated playback. This change in orientation is quite easy to do as long as the orientation of the slices is parallel to the x -, y -, or z -axes in the data set. If the depth direction is understood as the z -axis, then the x - and y -axes are the horizontal and vertical edges of the individual images. If the data are stored as discrete voxels, then accessing the data to form an image on planes parallel to these directions is just a matter of calculating the addresses of voxels using offsets to the start of each row and column in the array. This addressing can be done at real-time speeds if the data are held in memory, but is somewhat slower if the data are stored on a disk drive because the voxels that are adjacent along scan lines in the original slice images are stored contiguously on disk and can be read as a group in a single pass. When a different orientation is required the voxels must be located at widely separated places in the file, however and it takes time to move the head and wait for the disk to rotate.

Displaying an image in planes parallel to the x -, y -, and z -axes was introduced in Chapter 11. **Figure 12** shows another example of orthogonal slices. The images are MRIs of a human head. The views are generally described as transaxial (perpendicular to the subject's spine), sagittal (parallel

Figure 12. A few slices from a complete set of MRI head scan data. Images *a* through *c* show transaxial sections (3 from a set of 46), images *d* and *e* are coronal sections (2 from a set of 42), and *f* is a sagittal section (1 from a set of 30).






Figure 13. Comparison of two vertical slices through the 3D MRI data set from Figure 12:
 (a) slices extended vertically;
 (b) linear interpolation between slices.

to the spine and to the major axis of symmetry), and coronal (parallel to the spine and perpendicular to the “straight ahead” line of sight). Several individual sections are shown representing each orientation.

It is actually not too common to perform this kind of resectioning with MRI data (or most other kinds of medical images) because the spacing of the planes is usually greater than the resolution in the plane, and the result is a visible loss of resolution in one direction in the resectioned slices due to interpolation in the z direction. The alternative to interpolation is to extend the voxels in space; in most cases, this is even more distracting to the eye, as shown in **Figure 13**. Interpolation between planes of pixels can be done linearly, or using higher-order fits to more than two planes, or more than just the two pixels immediately above and below. But while interpolation produces a visually acceptable image, it can ignore real structure or create apparent structure. **Figure 14** shows an example of interpolation, which creates an impression of structure that is not actually present (and hides the real structure).

When data can be obtained with cubic voxels this interpolation is not a problem. In several of the figures later in this chapter, much greater loss of resolution in the z direction will be evident when plane images are reconstructed by sampling and interpolation of the original data. In the case of **Figure 12**, the MRI images were actually obtained with uniform resolution in each of the three directions.

Combining several views at once using orthogonal planes adds to the feeling of three-dimensionality of the data. **Figure 15** shows several examples of this using the same MRI head data based




Figure 14. Interpolation in a 3D array. In this example, only two images of the top and bottom of a woven fabric are used. The points between the two surfaces are linearly interpolated and bear no relationship to the actual 3D structure of the textile.




Figure 15. Several views of the MRI head data from **Figure 12** along section planes normal to the axes of the voxel array. The voxels were taken from the trans-axial slices, and so the resolution is poorer in the direction normal to the planes than in the planes.

on plane images obtained in the transaxial direction. The poorer resolution in the z direction is evident, but still the overall impression of 3D structure is quite good. These views can also be animated, by moving one (or several) of the planes through the data set while keeping the other orthogonal planes fixed to act as a visual reference.


Unfortunately, there is no good way to demonstrate this time-based animation in a print medium. Once upon a time, children's cartoon books used a "flip" mode with animation printed on a series of pages that the viewer could literally flip or riffle through at a fast enough rate to cause flicker-fusion in the eye and see motion. That form of animation takes a lot of pages and is really only good for very simple images such as cartoons. It is unlikely to appeal to the publishers of books and technical journals. All that can really be done here is to show a few of the still images from such a sequence and appeal to the reader's imagination to supply a necessarily weak impression of the effect of a live animation. Many online Web sites can show such animations as Quicktime® movies.

Figure 16 shows a series of images that can be used to show "moving pictures" of this kind. They are actually a portion of the series of images which Eadweard Muybridge recorded of a running horse by setting up a row of cameras that were tripped in order as the moving horse broke threads attached to the shutter releases. His purpose was to show that the horse's feet were not always in contact with the ground (in order to win a wager for Leland Stanford, Jr.). The individual still pictures show that. Once it was realized that such images could be viewed in sequence to recreate the impression of smooth motion, our modern motion picture industry (and ultimately television) became possible.

It is difficult to show motion using printed images in books. There is current interest in the use of videotape or compact disk for the distribution of technical papers, which will perhaps offer a medium that can use time as a third axis to substitute for a spatial axis and show 3D structure through motion. The possibilities will be mentioned again in connection with rotation and other time-based display methods.

Motion, or a sequence of images, is used to show multidimensional data in many cases. "Flipping" through a series of planes provides a crude method of showing data sets that occupy three spatial dimensions. Another effective animation shows a view of an entire 3D data set while varying the opacity of the voxels. Even for a data set that occupies two spatial dimensions, transitions between many kinds of information may be used effectively. **Figure 17** shows this multiplicity with weather data, showing temperature, wind velocity and other parameters displayed onto a map of the U.S. In general, displays that utilize a 2D map as an organizing basis for multidimensional data

Figure 16. A few of the series of historic photographs taken by Eadweard Muybridge to show the motion of a running horse. Viewed rapidly in succession, these create the illusion of continuous motion.



such as road networks, geological formations, and so on, called GIS, have many types of data and can only display a small fraction of it at any one time.

Of course, time itself is also a valid third dimension and the acquisition of a series of images in rapid succession to study changes in structure or composition with time can employ many of the same analytical and visualization tools as images covering three space dimensions. **Figure 18** shows a series of images recorded at video rate (30 frames per second) from a confocal light microscope. Such data sets can be assembled into a cube in which the z direction is time, and changes studied by sectioning this volume in planes along the z direction, or viewed volumetrically, or as a time sequence.





Figure 17. Weather data for the U.S. is a richly multidimensional data set tied to a geographic base.

Figure 18. A series of images of adult rat atrial myocytes loaded with the calcium indicator fluo-3. The images were recorded at video rate (30 frames per second) from a confocal light microscope (images courtesy of Dr. William T. Mason and Dr. John Hoyland, Department of Neurobiology, Babraham Institute, Cambridge, U.K.).




Arbitrary section planes

Restricting the section planes to those perpendicular to the x -, y -, or z -axes is obviously limiting for the viewer. It is done for convenience in accessing the voxels in storage and creating the display. If some arbitrary planar orientation is selected, the voxels must be found that lie closest to the plane. As for the case of image warping, stretching and rotating discussed in Chapter 3, these voxels will not generally lie exactly in the plane, nor will they have a regular grid-like spacing that lends itself to forming an image. **Figure 19** shows an example of a plane section through an array of cubic voxels, in which portions of various size and shape are revealed. These variations complicate displaying the voxel contents on the plane.

The available solutions are either to use the voxels that are closest to the section plane, plot them where they land, and spread them out to fill any gaps that develop, or to establish a grid of points in the section plane and then interpolate values from the nearest voxels, which may be up to eight in number. As for the case of rotation and stretching, these two solutions have different shortcomings. Using the nearest voxel preserves brightness values (or whatever the voxel value represents) but may distort boundaries and produce stair-stepping or aliasing. Interpolation makes the boundaries appear straight and smooth, but also smooths the brightness values by averaging so that

Figure 19. Intersection of a plane with an array of cubic voxels. The plane is viewed perpendicularly, showing the different areas and shapes of the intersection regions.



steps are blurred. It is also slower, because more voxels must be located in the array and the interpolation performed.

Producing an animation by continuously moving a section plane at an arbitrary orientation requires a significant amount of computation, even if it is simple calculation of addresses and linear interpolation of values. Instead of doing this calculation in real time, many systems instead create new images of each of the plane positions, store them, and then create the animation by playing them back. This procedure is fine for demonstrating something that the user has already found, but because of the time delay is not a particularly good tool for exploring the 3D data set to discover the unexpected.

The ultimate use of planar resectioning would be to change the location and orientation of the section plane dynamically in real time, allowing instant response to the scene. Chapter 1 pointed out that humans study things by turning them over, either in the hand or in the mind. This kind of "turning-over" is a natural way to study objects, but requires a fairly large computer memory (to hold all of the voxel data), a fairly speedy processor and display, and a user interface that provides the required number of degrees of freedom. Modern desktop computers are approaching this level of performance; **Figure 20** shows a typical display.



Figure 20. Presentation of serial sections, 3D reconstruction, and arbitrary plane slice on a desktop computer (courtesy Able Software).

Positioning an arbitrary plane can be done in two different ways, which, of course, produce the same results but feel quite different to the user. One is to move the plane with respect to a fixed 3D voxel array. This can be done, for example, by dragging the corners of the plane along the x , y , z -axes, or by positioning a center point and then tilting the plane around two axes through that point. A different method is to keep the section plane fixed perpendicular to the direction of view while allowing the data set to be rotated in space. Combined with the ability to shift the section plane toward or away from the viewer (or equivalently to shift the data set), this method allows exactly the same information to be obtained.

The principal difference between these approaches is that in the latter case the image is seen without perspective or foreshortening, so that size comparisons or measurements can be made. Obtaining such data may be important in some applications. On the other hand, keeping the data set fixed and moving the section plane seems to aid the user in maintaining orientation within the structure. There are as many different modes of human interaction with 3D data sets as there are programs, and as yet no general consensus on which are most useful has emerged.

Figure 21 shows the voxels revealed by an arbitrary section plane through the 3D data set from the MRI image series. The appearance of the voxels as a series of steps is rather distracting, so it is more common to show the value of the nearest voxels to the plane, or to interpolate among the voxels to obtain brightness values for each point on the plane as shown in **Figure 22**.

Also useful is the ability to make some of the pixels in section planes transparent, allowing the viewing of other planes behind the first, and making the 3D structure of the data more apparent. **Figure 23** shows an example of this for the spherical particle data shown in Chapter 11. The series of parallel slices are not contiguous planes of voxels, and the separation of the planes has been increased by a factor of two in the z direction. The voxels whose density falls below a threshold that roughly corresponds to that of the particles have been made transparent. This allows seeing voxels which are part of particles in other section planes behind the first.

Figure 24 shows a similar treatment for the MRI data set of the human head. The threshold for choosing which voxels to make transparent is more arbitrary than for the spheres, because void and low-density regions exist inside as well as outside the head. However, the overall impression of 3D structure is clearly enhanced by this treatment.




Figure 21. Sampling voxels along inclined planes in a 3D array. Showing the entire voxel is visually distracting and does not produce a smooth image

Figure 22. Smooth interpolation of image pixels on arbitrary planes positioned in a voxel array.






Figure 23. "Exploded" view of voxel layers in the tomographic reconstruction of spherical particles (Figures 54–55, Chapter 11). The low-density region surrounding the particles is shown as transparent.

Figure 24. View of several orthogonal slices of the MRI head data with transparency for the low-density regions in the planes.



The use of color

Assignment of pseudo-colors to grey-scale 2D images was discussed in earlier chapters. It sometimes permits distinguishing subtle variations that are imperceptible in brightness in the original; but, as noted before, it more often breaks up the overall continuity and gestalt of the image so that the image is more difficult to interpret. Of course, the same display tricks can be used with 3D sets of images, with the same consequences.

A subtle use of color or shading is to apply slightly different shading to different planar orientations shown in **Figures 15** and **22** to increase the impression of three-dimensionality. In this example, a grey-scale difference between the x , y , and z orientations is evident. Light tints of red, green and blue can be used for this as well.

It is more useful to employ different color scales to distinguish different structures, as was also demonstrated for 2D images. It requires separate processing or measurement operations to distinguish the different structures. When applied to 3D data sets, the colored scales assist in seeing the continuity from one image to another, while providing ranges of brightness values for each object.

One of the most common ways that multiple colors can be used to advantage in 3D image data is to code multiband data such as the elemental concentrations measured from the secondary ion mass spectrometer. This use of pseudo-color is analogous to similar coding in 2D images, very frequently used for x-ray maps from the SEM and (of course) for remotely sensed satellite images, in which the colors may either be "true" or used to represent colors beyond the range of human vision, particularly infrared. Of course, the other tools for working with multiband images in 2D, such as calculating ratios, can also be applied in 3D.

Figure 25 shows example images for the SIMS depth imaging of elements implanted in a silicon wafer. Comparing the spatial location of the different elements is made easier by superimposing the separate images and using colors to distinguish the elements. This is done by assigning red, green and blue to each of three elements and then combining the image planes, with the result shown in **Figure 26**.

Volumetric display

Sectioning the data set, even if some regions are made transparent, obscures much of the voxel array. Only the selected planes are seen and much of the information in the data set is not used. It has been pointed out before that the topological properties of 3D structures are not revealed on





Figure 25. Views of plane sections for the elements aluminum, boron, and oxygen in a silicon wafer, imaged by secondary ion mass spectrometry. **Figure 26** shows a color image of all three elements on another orthogonal set of planes.

Figure 26. Color coding of elemental intensity from SIMS images in **Figure 25**. The multiband 3D data set is sectioned on *x*, *y*, and *z* planes, and the 256-level (8-bit) brightness scale for the elements Aluminum, Boron, and Oxygen are assigned to red, green, and blue, respectively.



section planes. The volumetric display method shows all of the 3D voxel information. For simple structures displaying everything can be an advantage, while for very complex ones the overlapping features and boundaries can become confusing.

A volumetric display is produced by ray tracing. In the simplest model used, a uniform, extended light source is placed behind the voxel array. For each parallel straight line ray from the light source to a point on the display screen, the density value of each voxel that lies along the path is used to calculate a reduction in the light intensity following the usual absorption rule:

$$\frac{I}{I_0} = e^{-\sum \rho} \quad (3)$$

Performing this calculation for rays reaching each point on the display generates an image. The total contrast range may be adjusted to the range of the display by introducing an arbitrary scaling constant. This scaling can be important because the calculated intensities may be quite small for large voxel arrays.

Notice that this model assumes that the voxel values actually correspond to density, or to some other property that may be adequately modeled by the absorption of light. The images shown in **Figures 54** and **56a** in **Chapter 11** correspond closely to this model, since they show the projected view through a specimen using x-rays, although the geometry is cone beam rather than parallel projection. In fact, x-ray tomographic reconstruction, discussed in the preceding chapter, proceeds from such views to a calculation of the voxel array. Having the array of voxel values then permits generating many kinds of displays to examine the data. It seems counter-productive to calculate the projection view again, and indeed in such a view as shown in those figures, the ability to distinguish the individual features and see their relationship is poor.

One of the advantages of this mode is that the direction of view can be changed rather easily. For each, it is necessary to calculate the addresses of voxels that lie along each ray. When the view direction is not parallel to one of the axes, this addressing can be done efficiently using sine/cosine values or a generalization of the Bresenham line drawing algorithm. Also in this case, an improved display quality is obtained by calculating the length of the line segment along each ray through each pixel. The absorption rule then becomes

$$\frac{I}{I_0} = e^{-\sum \rho l} \quad (4)$$

This method is far short of a complete ray tracing, although it is sometimes described as one. In a true ray-traced image, refraction and reflection of the light is included along with the absorption.



Figure 27. Reconstruction of chromosomes in a dividing cell from CSLM 3D data. The chromosomes are opaque and the matrix around them transparent, and shadows have been ray cast on the rear plane to enhance the 3D appearance.

Figure 27 shows an example in which the inclusion of shadows greatly increases the 3D impression. More complex shading, in which features cast shadows on themselves and each other, requires calculations that are generally too time-consuming for routine use in this application. With the simple absorption-only method, it is possible to achieve display speeds capable of rotating the array (changing the view direction) interactively with desktop computers.

Of course, it is always possible to generate and save a series of projection images that can then be played back as an animation or movie. These are primarily useful for communicating some information already known to another viewer, however, while interactive displays may assist in discovering the structural relationships in the first place.

These types of volumetric displays are often isometric rather than perspective-corrected. In other words, the dimension of a voxel or feature does not change with distance. This is equivalent to looking at the scene through a long focal length lens, and given the inherent strangeness of data in most 3D image sets does not generally cause significant additional discomfort to viewers. True perspective correction requires that x,y dimensions on the screen be adjusted for depth. Particularly for rotated views, perspective adds a significant amount of computation.

Figure 28 shows such a view through a joint in the leg of a head louse. The original series of images were obtained with a transmission confocal light microscope, with nearly cubic voxels (the spacing between sequential focal planes was $0.2\text{ }\mu\text{m}$ in depth). The individual muscle fibers are visible but overlapped. Shifting the stack to approximate rotation as described below gives the viewer the ability to distinguish the various muscle groups. Again, a time sequence (hard to show in print media) is used as a third dimension to display 3D data as the planes are shifted, and since no complex arithmetic is needed to run this sequence it is practical to create such animations in a small computer.

A much more limited approach to volumetric presentation is to show one portion of the image as a surface rendered model, as discussed below, but to show the surrounding tissue as a transparent

Figure 28. Volumetric projection image through a stack of 60 CSLM images of a joint in the leg of a head louse. Each plane is displaced by one voxel dimension to produce a view at 45° .





Figure 29. Viewing the surface reconstruction from serial sections through a knee joint with the surrounding tissue shown as a transparent volume with no internal detail.




volume. Usually no detail is preserved in this volume, and it is present only to provide a point of reference for the structure of interest, as shown in **Figure 29**.

Very rapid generation of projection images is possible if the addressing can be simplified and the variation in distance through different voxels can be ignored. **Figure 30** shows an approximation that facilitates these changes. Each plane of voxels is shifted laterally by a small amount, which may be an integer number of voxel spaces making the address calculation particularly simple, but in any case requires no more than simple 2D interpolation. The planes remain normal to the view direction, so that all distances through pixels are the same. This kind of shifting can give the impression of rotation for small angles. Beyond about 30° , the distortion of the 3D structure due to stretching may become visually objectionable; however, this method provides a fast way to cause some relative displacement of features as a function of depth to better understand the structure.

Stereo viewing

In many of the images in this chapter, two adjacent images in a rotation or pseudo-rotation sequence can be viewed as a stereo pair. For some readers looking at them will require an inexpensive viewer which allows one to focus on the separate images while keeping the eyes looking straight ahead (which the brain expects to correspond to objects at a great distance). Other readers may have mastered the trick of fusing such printed stereo views without assistance. Some,

Figure 30. Schematic diagram of shifting image planes laterally to create illusion of rotation, or to produce stereo-pair images for viewing.



unfortunately, will not be able to see them at all. A significant portion of the population seems not to actually use stereo vision, due for instance to uncorrected amblyopia ("lazy eye") in childhood.

Stereo views are so useful to a reasonable fraction of people that it may be useful to display them directly on the viewing screen. Of course, with a large screen, it is possible to draw the two views side by side. **Figures 31** and **32** show examples of stereo pair presentation of 3D images using both the volumetric display method discussed above and the surface-rendered method discussed below, for the same specimen (derived from confocal light microscope images of a sea urchin embryo). **Figure 33** shows another stereo pair presentation of skeletonized data obtained from neurons imaged in the confocal light microscope.

A more direct stereo display method uses color planes in the display for the left- and right-eye views. For instance, **Figure 34** shows a stereo pair of blood vessels in the skin of a hamster, imaged live using a confocal light microscope in fluorescence mode. The 3D reconstruction method used pseudo-rotation by shifting of the focal section planes with the emission rules discussed




Figure 31. Stereo pair presentation of Feulgen-stained DNA in a sea urchin embryo. The cells form a hollow sphere, evident in the stereo images, with some cells in the act of cell division. This image shows a volumetric image that is "ray cast" or "ray traced," using an emission model, with different offsets for the individual optical sections to produce a stereo effect. See also **Figure 32** for a surface image of the same data (From Summers, R.G. et al. (1991). J. Electron Microscope Tech. 18:24–30. With permission.)




Figure 32. Stereo view of the data set from **Figure 31**, but surface is rendered and color coded. The surface image shows contours within each section which render the surfaces of the chromosomal masses but obscure any internal detail or structures to the rear. Contour lines for the embryo are also shown.

Figure 33. Stereo images of skeletonized lines (manually entered from serial section data) from two neurons in hippocampus of 10-day old rat, showing branching and 3-D relationships (From Turner, J.N. et al. (1991). J. Electron Microscope Tech. 8:11–23. With permission.)




Figure 34. Stereo view of multiple focal plane images from a confocal light microscope showing light emitted from fluorescent dye injected into the vasculature of a hamster and viewed live in the skin. This image is also shown in **Figure 35** (images courtesy Chris Russ, University of Texas, Austin TX).

below. Combining these images using red and green to display two views of the same 3D data set is shown in **Figure 35**. The images are overlapped, and the viewer (equipped with glasses having appropriate red and cyan filters) can easily look at the combined images. Of course, this method cannot be used for color images, as discussed in the following paragraphs.

These stereo views were constructed from multiple sections by projecting or ray tracing through a stack of images as shown in **Figure 30**. The individual sections were obtained by confocal microscopy, which, due to its very shallow depth of field, can be used to obtain a continuous set of voxel planes. Other techniques that produce continuous arrays of thin sections can also be used, of course; but it is even possible to perform this type of reconstruction using a conventional optical microscope, despite the blurring of the images due to the large depth of field and the effect of light passing through the specimen above and below the image planes.

Removing artefacts such as blur from images by processing in frequency space was discussed in Chapter 5. For three dimensions the procedure is identical except that the point spread function and the Fourier transform are three dimensional. **Figure 36** shows an example of using Wiener inverse filtering to accomplish this deblurring in the creation of a stereo pair. The images were reconstructed from 90 optical sections spaced through a 50- μ m thick section, which is much closer than the depth of field of the optics. The use of the inverse filter removes most of the artefacts from the images and produces a clear stereo pair image (Lin et al., 1994). Another approach to the same kind of sharpening is to apply an iterative procedure that uses neighboring images to estimate the





Figure 35. Stereo pair of the same image pair shown in Figure 34, using red and cyan for the different eye views. This allows viewing the image with normal eye vergence, using glasses (red lens on left eye, green or blue on right).

Figure 36. Sharpening of focal sections by Wiener filtering. The stereo-pair images are produced by ray tracing two projections at angles of about $\pm 2.4^\circ$ through 90 serial optical sections in a 50 μm thick section. The sample is *Spirogyra*. (a) original images; (b) Wiener inverse filtered (From Lin, W. et al. (1994). J. Comput. Assisted Microsc. 6(3):113–128. With permission.)



artefacts in each plane. This Van Cittert filter requires using neighbor planes on each side out to about twice the dimension of the point spread function, which in this case is the depth of optical focus (Jain, 1989).

For projection of stereo pair images, it is possible to use two slide projectors equipped with polarizing filters that orient the light polarization at right angles (usually 45° to the right and left of vertical). Viewers wearing polarized glasses can then see the stereo effect, and color can still be used. This display method requires special specular projection screens that reflect the light without losing the polarization, and works best for viewers in line with the center of the screen. However, it has become a rather popular method of displaying 3D data. Of course, it is not as practical for interactive

exploration of a data set since photographic slides must be made first. Polarization can also be used with live computer displays, as discussed in the next paragraph.

The difference in viewing angle for the two images can be adjusted somewhat arbitrarily to control the visual impression of depth. The angle can be made to correspond to the typical vergence angle of human vision for normal viewing. Using a typical interocular distance of 7.5 cm, and a viewing distance of 1 meter, the actual vergence angle is 4.3° . For closer viewing, larger angles are appropriate. The judgment of depth thus depends on our brain's interpretation of the viewing distance, which is based on the focus distance to the image in combination with the vergence angle of the eyes in their sockets. If the angle is varied, the impression of depth can be adjusted to expand or compress the z dimension, as shown in **Figure 37**.




Figure 37. Stereo pair images of a generated structure with varying angles to control the apparent depth of the structure in the z direction: (a) $\pm 1^\circ$; (b) $\pm 3^\circ$; (c) $\pm 5^\circ$. Visual fusion of the images become increasingly difficult with the larger angles.

Special display hardware

Other specialized display hardware for 3D image analysis may be useful in some cases. Holography offers the promise of realistic 3D display that can be viewed from different directions (Blackie et al., 1987). Attempts to generate such displays by calculating the holograms have been experimentally successful, although they are still too slow for interactive use (and far from the “stand alone” 3D presentation of Princess Leia in *Star Wars*). At present, the best holograms are displayed using coherent light from a laser and high resolution film. In order to produce live displays from a computer, multiple screens such as the LCDs used for conference room projection can be used in place of the film; however, the resolution and control of the light modulation (relative intensity) is not really adequate.


Another custom approach is called a varifocal mirror (Fuchs et al., 1982). Each plane of voxels in the 3D array is drawn one at a time on the display CRT. The screen is not viewed directly by the user, but is reflected from a mirror. The mirror is mounted on a speaker voice coil so that it can be moved. As each different plane of voxels is drawn, the mirror is displaced slightly as shown schematically in **Figure 38**. This movement changes the distance from the viewer's eye to the screen, and gives the impression of depth. In order to achieve high drawing speeds (so that the entire set of planes can be redrawn at least 30 times per second), this technique is usually restricted to simple outline drawings rather than the entire voxel data set. The successive outlines are perceived as surfaces in three dimensions.

The mirror technique also suffers from a slowdown if the planes are not normal to the x -, y -, and z -axes so that trigonometric calculations are required to access the data. The alternative approach is to continue to draw outlines in the major orthogonal planes, but to move the mirror during the drawing of each outline to correspond to the tilting of the plane. This requires a much higher and more complex frequency response from the mechanical devices used to control the mirror.

Another, more recent development for real-time viewing of 3D computer graphics displays uses stereo images. The computer calculates two display images for slightly different orientations of the data set or viewing positions. These are displayed alternately using high-speed display hardware which typically shows 120 images per second. Special hardware is then used to allow each eye to see only the correct images, at a rate of 60 times per second, fast enough to eliminate flicker (the minimum rate for flicker fusion above which we see continuous images rather than discrete ones is usually at least 16 frames per second; commercial moving pictures typically use 24 frames, and television uses 25 for European and 30 in U.S. systems).

The visual switching may be done by installing a liquid crystal device on the display monitor that can rapidly switch the polarization direction of the transmitted light, so that viewers can watch through glasses containing polarizing film. A second approach is to wear special glasses containing active liquid crystal devices that can rapidly turn clear or opaque. Synchronizing pulses from the computer cause the glasses to switch as the images are displayed, so that each eye sees the proper view.

Figure 38. Diagram of the operation of a varifocal mirror to show depth in images. The speaker voice coil rapidly varies the position of the mirror, which changes the distance from the viewer's eye to the cathode ray tube display as it draws information from different depths in the data set.



Such devices have been used primarily for graphics design, in which substantial computer resources are used to model 3D objects, generate rendered surface views, and allow the user to freely rotate and zoom. With the number of disciplines interested in using 3D computer graphics, it seems assured that new hardware (and the required corresponding software) will continue to evolve for this purpose. The economic breakthrough that would impact scientific uses will probably come when game console manufacturers decide to deliver 3D displays.

These various display tools can be adapted to the display of 3D image data and used at scales ranging from nanometers (electron and ion microscopy) to kilometers (seismic exploration). It is possible to imagine using other senses than visual to deal with multiband data (e.g., sound which changes pitch and volume to reveal density and composition as you move a cursor over the display of a voxel array), but as emphasized in Chapter 1, we are overwhelmingly visual creatures who expect to get information through images. Existing technologies such as force-feedback gloves that would allow the user to "reach into" structures and feel or manipulate them, and virtual reality "caves" using lightweight headsets with stereo displays that sense the user's head position and generate appropriate images that allow one to wander through the microstructure, are still much too costly for routine use. Because vision is our primary sense, however, it appears likely that most methods will be primarily visual, and the researcher who is color-blind or cannot see stereo will remain disadvantaged.

Also, no consensus has been reached on the best input and control devices for complex computer graphics displays. Rotating or shifting the viewpoint in response to horizontal or vertical motion of the now ubiquitous mouse gives a rather crude control. For moving points, lines, and planes in 3D, some more flexible device will be required. Simply locating a specific voxel location in the array can be done in several different ways. One is to use an x,y input device (mouse, trackball, etc.) for two axes and periodically shift to a different control, which may be a scroll bar on the screen, to adjust the distance along the third axis. Another is to move two cursors, one on an $x-y$ projection and one on an $x-z$ projection, for example. Such approaches usually feel rather clumsy because of the need to move back and forth between two areas of the screen and two modalities of interaction. Appropriate color-coding of the cursor to report depth can help.

Three-axis joysticks and sonic digitizers that allow the user to point in space (multiple sensors triangulate the position) exist but are hardly standard. The dataglove, an instrumented glove or framework that fits on the hand and reports all motion of joints to the computer, has been used to move molecules around each other to study enzyme action. Supplemented by force feedback, this method gives the researcher rich information about the ways that molecules can best fit together.

At the other extreme, it is important to remember the history of 3D visualization (Cookson, 1994), which began with physical models constructed of wood, plastic or plaster. Building such a model from a series of section planes was very difficult and time-consuming, and still could not reveal all of the internal detail. Computer modeling has progressed from simple outlines to hidden line removal, surface construction, shading and rendering, and full volumetric or ray-traced methods.

Ray tracing

The example of volumetric display shown above performed a simplified ray tracing to sum the density values of voxels and calculate a brightness based on light being absorbed as it propagated from back to front through the 3D data set. Although this model does correspond to some imaging situations, such as the transmission light or electron microscope and tomography, different rules are appropriate in many other situations.

In the process of traversing a voxel array, following a particular line of sight that will end in one pixel of a ray-traced image, the variables that are available include:

1. The brightness and perhaps color of the original light source placed behind the array, and whether it is an extended source (producing parallel rays of light) or a point source (producing a cone beam of light). This illumination source will control the contribution that transmission makes to the final image.
2. The location of the first and last voxels with a density above some arbitrary threshold taken to represent transparency. These voxels will define surfaces that can be rendered using reflected light. Additional rules for surface reflectivity, the location, brightness, and color of the light sources, and so forth, must be added.
3. The location of the maximum or minimum values or large gradients, which may define the location of some internal surface for rendering.
4. The rule for combining voxel values along a path. This may be multiplication of fractional values, which models simple absorption according to Beer's law for photons provided that the voxel values are linear absorption values. In some cases density is proportional to attenuation, so this rule can produce interpretable images. There are other convolution rules available as well, including linear summation and retention of maximum or minimum values. Although these may also correspond to some physical situations, their greatest value is that they produce images which can delineate internal structure.
5. The relationship between the voxel values and the intensity (and perhaps color) of light originating in each voxel, which represents fluorescence or other emission processes.

The combining rules mentioned briefly in part 4 of the above list correspond to the various image processing tools described in Chapters 3 and 4 for combining pixels from two or more images. They include arithmetic (multiplication, addition), rank ordering (minimum or maximum value), and Boolean logic. It is also possible to include lateral scattering so that point sources of light spread or blur as they pass through the voxel array, or even to combine several modes. This approach to realism through computation is rarely justified because the measured voxel values are not generally physically related to light transmission or scattering.

A software package for 3D visualization may make any or all of these parameters accessible to the user, along with others. For example, control of the surface reflectivity and roughness, and the location of the incident light source(s), affects the appearance of rendered surfaces. In performing a convolution of transmitted light along ray paths from a light source behind the voxel array, the relationship between the voxel values and the absorption of the light is another parameter that offers control. By varying the relationship between voxel value and opacity (linear attenuation coefficient), or selectively introducing color, it is possible to make some structures appear or to remove them, allowing others to be seen.

Figure 39 shows an example of this presentation. The data set is the same as used above; the voxel values come from the MRI measurement technique and approximately correspond to the amount of water present. Not enough information is given to fully describe the structures actually present in a human head, so there is no "correct" relationship between voxel value and light absorption for the volumetric rendering. Using different, arbitrary curves, it is possible to selectively view the outer skin, bone structure, or brain.

Using color permits even more distinctions to be made. **Figure 40** shows images of a hog heart reconstructed using different relationships for opacity vs. voxel value that emphasize the heart muscle or the blood vessels. In this example, each voxel is assumed both to absorb the light from the source placed behind the voxel array and to contribute its own light along the ray in proportion to its value and with color taken from an arbitrary table. The result allows structures with different measured values to appear in different colors.

Of course, with only a single value for each voxel it is not possible to model absorption and emission separately. By performing dual-energy tomography in which the average atomic number and

Figure 39. Volumetric imaging of the MRI head data from **Figure 12**. Varying the relationship between voxel values and the opacity used to absorb light transmitted along each ray through the structure allows selection of which structures are revealed.




Figure 40. Volumetric rendering of MRI data. Specimen is a hog heart (data courtesy B. Knosp, R. Frank, M. Marcus, R. Weiss, University of Iowa Image Analysis Facility and Department of Internal Medicine). Changing the arbitrary relationship between voxel value and display opacity for the voxels allows selectively showing the heart muscle or blood vessels.

average density are both determined, or multi-energy tomography in which the concentration of various elements in each voxel is measured, or by using the T1,T2 relaxation time signals from MRI, such techniques become possible. They represent straightforward implementation of several of the multiband and color imaging methods discussed in earlier chapters, but are not yet common as neither the imaging instrumentation nor the computer routines are yet widely available. It is consequently usually necessary to adopt some arbitrary relationship between the single measured set of voxel values and the displayed rendering that corresponds to the major voxel property measured by the original imaging process.

For example, in fluorescence light microscopy, or x-ray images from the SEM, or ion microscopy, the voxel value is a measure of emitted brightness that is generally proportional to elemental concentration. These 3D data sets can also be shown volumetrically by a simplified ray tracing.

Instead of absorbing light from an external light source, the rule is to sum the voxel values as brightnesses along each path.

Figure 34 showed an application using the fluorescence confocal light microscope. A dye was injected into the blood vessel of a hamster, which was excited by the incident light from the microscope. The emitted light was collected to form a series of 2D images at different focal depths, and these were then arranged in a stack to produce a 3D data set. In this case the spacing between the planes is much greater than the resolution within each image plane. Sliding the image stack laterally as discussed above produces an approximation of rotation and impression of depth. The brightness values for each voxel are then summed along vertical columns to produce each image.

This emission model is very easy to calculate but does not take into account any possible absorption of the emitted light intensity by other voxels along the ray path. Generally, simple 3D data sets have only one piece of data per voxel and there is no separate information on density and emission brightness, so no such correction is even possible. Sometimes a simple reduction in intensity in proportion to the total number of voxels traversed (known as a "distance fade") may be used to approximate this absorption effect. Usually, it is assumed that the emission intensity is sufficiently high and the structure sufficiently transparent that no such correction is needed, or that it would not change the interpretation of the structure, which is in any case qualitative rather than quantitative.

When multiband data are available, as for instance in the SIMS data set used in **Figures 25** and **26**, emission rules can be used with the assignment of different colors (at least up to three) to different signals. **Figures 41** and **42** show a volumetric view of these data using emission rules, presented as a stereo pair. The monochrome figure shows a single element (boron) while in the color image, multiple elements are combined. The use of color in the images forces the use of two

Figure 41: Stereo-pair display of emission rule volumetric images of Boron concentration in a silicon wafer, imaged by a secondary ion mass spectrometer.





Figure 42. Stereo pair of volumetric images from the SIMS data set. Colors assigned to the different elements are Red = Aluminum, Green = Boron, and Blue = Oxygen, with 256 brightness levels of each. The use of color for elemental information precludes color coding of views for stereo viewing. The density of information in this display is so great that fusing the images to see the depth is quite difficult.

Figure 43. Surface rendered display of the same data as in **Figure 42**. The internal surface of the boron-rich region is determined by arbitrary thresholding, and then rendered using an arbitrarily placed light source.



side-by-side images for viewing, and the density of information in the images that results from overlaying the 8-bit (256 grey-level) values from each element at every point makes it quite difficult to fuse these images for satisfactory stereo viewing.

Using the same data set, it is possible to use the location of the frontmost voxel along each line of sight whose value is above an arbitrary threshold to define the location of an internal boundary. If this resulting surface is then rendered as a solid surface with incident reflected light, another representation of the data is obtained as shown in **Figure 43**.

An internal surface can be displayed in several different ways. **Figure 44** shows several views of the spiral cochlear structure from the ear of a bat (Keating, 1993). Slices through the 3D voxel array do not show that the structure is connected. Thresholding to show all of the voxels within the structure reveals the topology of the spiral. A series of section images can be shown as either wire frame outlines, or with the surface rendered. The outlines are more convenient for rotating to view the structure from different directions.

When color coding of structures, volumetric rendering, and selective transparency are combined with the display of internal surfaces (with shading as discussed in the next section), the results can be quite striking. **Figure 45** shows an example, the brain of a fruit fly (*Drosophila melanogaster*), using data recorded with a confocal laser scanning microscope. The 3D reconstruction shows surfaces rendered as partially transparent in a transparent medium, to reveal the surfaces of other structures behind them. Although effective at communicating specific information to the viewer, these displays require prior decisions and selection and are not so useful for exploring unknown data sets, as noted previously.

Reflection

Another important imaging modality is reflection. The conventional CSLM, seismic reflection mapping, acoustic microscopy, ultrasound imaging, and so forth, acquire a 3D image set whose voxel values record the reflection of the signal from internal locations. **Figure 46** shows the use of ultrasound to show the surface of a fetus in utero. In most of these technologies, a voxel array is generated and used to locate boundaries where strong reflections occur, within a matrix that may be otherwise opaque. In the CSLM, the matrix is transparent (either air or liquid) and the strongest reflection at each x,y point is where the specimen surface is in focus. This means that recording a set of images as a 3D data array makes it possible to locate the surface in three dimensions. Most systems use the processing methods discussed in Chapter 4 to find the brightest value at each pixel, and thus construct the surface range image.

Figure 44. Cochlear structure from the ear of a bat (Keating, 1993):

- (a) arbitrary planar slices through the voxel array;
- (b) array of parallel slices through the array;
- (c) surface revealed by thresholding the voxels;
- (d) outlines delineating the structure in all voxel planes;
- (e) surface reconstructed from the outlines in d;
- (f) the same structure as shown in e, rotated to a different point of view.





Figure 45. Reconstruction of a fruit fly brain with transparent surfaces. Data courtesy of K. Rein (Department of Genetics, University Wuerzburg), visualization by M. Zoeckler (ZIB) using the Amira software (www.amiravis.com).



One way to generate such a display is to go down columns of the data array (just as in volumetric ray tracing), looking for the maximum voxel value, or the first value to exceed some threshold. Keeping only that value produces an image of the entire surface in focus, as was shown in Chapter 4. The same methods for rotating or shifting the data array to alter the viewing direction can be used; therefore, it is also possible to find the maxima or surface points along any viewing direction and display the surface as an animation sequence or to construct a stereo pair. In principle, fitting a curve based on the known depth-of-field characteristics of the optics to the brightness values along a vertical column of voxels can locate the surface with sub-voxel accuracy. **Figure 47**

Figure 46. The face of a 26-week-old fetus, imaged by ultrasound.




Figure 47. Depth

measurement using the confocal scanning light microscope:

(a) extended focus image of an alumina fracture surface obtained by keeping the brightest value from many focal planes at each pixel location;

(b) range image obtained by assigning a grey scale value to each pixel according to the focal plane image in which the brightest (in focus) reflectance is measured;

(c) elevation profile along the traverse line shown in figure b.

shows the alumina fracture (from **Chapter 4, Figure 80**) both as an extended-focus image and as a range image (in which pixel brightness is proportional to elevation). Several presentation modes for range images are available to assist in visualizing the 3D shape of the surface. One, shown in the figure, is simply to plot the brightness profile along any line across the image, which gives the elevation profile directly.

Figure 48 shows several of the presentation modes for range images (Chapter 13 goes into more detail). The specimen is a microelectronics chip imaged by reflection CSLM, so both an extended-focus image and a range image can be obtained from the series of focal plane sections. From the range image data, plotting contour maps, grid or mesh plots, or rendered displays is a straightforward exercise in computer graphics.

One of the classic ways to show surface elevation is a contour map (**Figure 49**), in which isoelevation lines are drawn, usually at uniform increments of altitude. These lines are of course

Figure 48. Presentation modes for surface information from the CSLM (specimen is a microelectronics chip):

- (a) the in-focus image of the surface reflectance obtained by keeping the brightest value at each pixel address from all of the multiple focal plane images;
- (b) the elevation or range image produced by grey-scale encoding the depth at which the brightest pixel was measured for each pixel address;
- (c) contour map of the surface elevation with color coding for the height values;
- (d) perspective-corrected rendering of the surface with grid lines and pseudo-color;
- (e) the same image as d with realistic shading;
- (f) the rendering from e from two points of view, combined as a stereo pair.





Figure 49. Fragment of a conventional topographic map (Phantom Ranch in the Grand Canyon) showing isoelevation contour lines.



Figure 50. Isometric view of elevation data for the alumina fracture surface shown in **Figure 47**. A rapidly varying color palette is used to reveal small changes.



continuous and closed. This is the way topographic maps are drawn, and the same methods are useful at any scale. The contour map reduces the pixel data from the original range image to boundary representation, so the method for forming the boundaries is the same as discussed in Chapter 5 for segmentation. The lines may be labeled or color-coded to assist in distinguishing elevations.

A shortcut way to draw contour lines on a range image is to present the data as a shaded isometric view, as in **Figure 50**, which shows the elevation data for the alumina fracture surface of **Figure 47**. In this image, a 3D representation (without perspective) is used to draw a vertical line for each pixel in the range image to a height proportional to the value. The image is also shaded so that each point has its grey-scale value. Replacing the grey-scale values with a pseudo-color table allows communication of the elevations in a particularly easy-to-interpret way, and in fact many topographic maps use similar methods to show elevations.

Constructing a contour map can be either simple or complex. The simplest method merely locates those pixels that have neighbors that are above and below the threshold level. As shown in **Figure 51**, however, this produces a very coarse approximation. Interpolating between pixel addresses produces a much better map, as shown in the figure.

Figure 52 shows an example that looks broadly similar to **Figure 51**, but represents data at a very different scale. This is a three-dimensional view of Ishtar Terra on Venus. The data come from the spacecraft Magellan's side-looking mapping radar. This synthetic-aperture radar bounces 12.5-cm-wavelength radar waves off the surface, using the echo time delay for range, and the Doppler shift to collect signals from points ahead of and behind the direct line of sight. The 2D images obtained by processing the signals are similar in appearance to aerial photographs.

Rendering of a surface defined by a range image produces a realistic image of surface appearance, as compared to grid or isometric contour map displays that are more abstract and more difficult for visual interpretation, as shown in **Figure 53**; however, the quantitative interpretation of the surface data is more readily accessible in the range image. Also, it is difficult to select realistic surface colors and textures to be applied. It is possible to apply brightness values to an isometric display of range data that come from another image of the same area, such as the original reflectivity or texture information. When multiband images are recorded, this combination is particularly effective.





Figure 51. Surface of a metal fracture:
(a) range image (grey scale proportional to elevation);
(b) shaded grid;
(c) contour map produced by selecting pixels with neighbors above and below each of five thresholds;
(d) contour map produced by linear interpolation.

Figure 52. Reconstructed surface image of Ishtar Terra on Venus, looking northeast across Lakshmi Panum toward Maxwell Montes with an exaggerated vertical scale (From Saunders, S. (1991). Eng. Sci. Spring 1991:15-27. With permission.)



The surfaces discussed in the preceding section are in fact external physical surfaces of the specimen. Internal surfaces may be defined as boundaries between distinct regions, or in more subtle ways. For instance, **Figure 54** shows the data from the SIMS example used above, in which the depth of the voxel having the maximum concentration of silicon at any location is shown. This surface isolation provides a visualization of the shape of the implanted region. The image is shown as a shaded isometric display, as discussed previously.

Surfaces

Surfaces to be examined may be either physical surfaces revealed directly by reflection of light, electrons or sound waves, or surfaces internal to the sample and revealed only indirectly after the entire 3D data set has been acquired. The use of computer graphics to display them is closely related to other graphics display modes used in CAD, for example; however, the typical CAD object

Figure 53. Display modes for surface information:

- (a) range image;
- (b) grid mesh;
- (c) isometric view;
- (d) rendered terrain.




Figure 54. Isometric display of elevation data within a volume. The height of the surface having the maximum concentration on Si in the SIMS voxel array.

has only a few numbers to describe it, such as the coordinates of vertices. Generating the interpolated surfaces and calculating the local orientation and hence the brightness of the image at many points requires a significant amount of computing.

In contrast, the image data discussed here is typically a complete 3D data set, or at least a complete 2D range image derived from the 3D set. Consequently, there is elevation data at every pixel location in the display image, which allows for an extremely rapid image generation. Rendering the surface images shown here took only a few seconds on a desktop computer. Doing the same for a typical CAD object would take longer.




Figure 55. Presentation modes for the surface elevation data from an optical interferometer (specimen is the machined surface of nickel):


- (a) original image, in which the grey-scale brightness encodes height;
- (b) a contour map with grey-scale shading to indicate the height of lines;
- (c) isometric view with superimposed grey scale values from image a;
- (d) the surface data rendered as it would appear with a diffuse material;
- (e) the surface data rendered as it would appear with a specular material.

Some instruments produce range images directly. Large-scale examples include radar mapping, elevation measurement from stereo pair calculations, and sonar depth ranging. At a finer scale, a standard tool for measuring precision machined surfaces is interferometry, which produces images as shown in **Figure 55**. The brightness is a direct measure of elevation, and the image can be comprehended more easily with appropriate rendering. Notice that the lens artefact (the faint ring structure at the left side of the image) is not true elevation data and when presented as such looks quite strange.

Displays of surface images (more formally of range images, because real surfaces may be complex and multivalued, but range images are well behaved and single valued) can use any of the techniques described above. These include wire mesh or line profile displays, contour maps, and shaded isometric displays, all described in more detail in Chapter 13. These all involve a certain level of abstraction.

A simple set of line profiles gives an impression of surface elevation and requires no computation, although the need to space the lines apart loses some detail. Consequently, it is sometimes used as a direct display mode on instruments such as the SEM or scanning tunneling microscopy (STM). Unfortunately, the signal that is displayed in this way may not actually be the elevation, and in this case the pseudo-topographic display can be quite misleading. Adding grid or mesh lines in both directions requires additional computation, but also increases the effective spacing and decreases the lateral resolution of the display.

Figure 56. Lambertian rules for light scattering from surfaces with different specularities. The vectors are N (surface normal), L (light source), R (reflection), and V (viewing direction). the k factors are the diffuse and specular reflection coefficients. The I values are the intensities of the principal and ambient light sources, and b is a constant describing the breadth of the specular reflection, which depends on the fine-scale surface roughness.



Generating an image of a surface that approximates the appearance of a real, physical surface is known generically as rendering and requires more computational effort. The physical rules that govern the way real surfaces look are simple and are summarized in **Figure 56**. The important variables are the intensity and location of the light source and the location of the viewer. Both are usually given in terms of the angles between the normal vector of the surface and the vectors to the source and viewer. The absolute reflectivity of the surface (or albedo) must be known; if this varies with wavelength we say that the surface is colored because some colors will be reflected more than others.

Finally, the local roughness of the surface controls the degree of variation in the angle of reflection. A very narrow angle for this spread corresponds to a smooth surface that reflects specularly. A broader angle corresponds to a more diffuse reflection. One of the very common tricks in graphic arts, which can be seen any evening in television advertising, is the addition of bright specular reflections to objects to make them appear metallic and hopefully more interesting.

For a typical surface defined by a few points, as in CAD drawings, the surface is broken into facets, often triangular, and the orientation of each facet is calculated with respect to the viewer and light source. The reflected light intensity is then calculated, and the result plotted on the display screen or other output device to build the image. This would seem to be a rather fast process with only a small number of facets, but the problem is that such images do not look natural. The large, flat facets and the abrupt angles between them do not correspond to the continuous surfaces we encounter in most real objects.

Shading the brightness values between facets (Gouraud shading) can eliminate these abrupt edges and improve the appearance of the image, but requires much interpolation. Better smoothing can be achieved (particularly when there are specular reflections present) by interpolating the angles between the centers of the facets rather than simply the mean brightness values. The interpolated angles are used to generate local brightness values, which vary nonlinearly with angle and hence position. This Phong shading is even more computer-intensive.

For continuous pixel images, each set of three pixels can be considered to define a triangular facet as shown schematically in **Figure 57**. The difference in value (elevation) of the neighboring pixels gives the angles of the local surface normal directly. A precalculated lookup table (LUT) of the image brightness values for a given light source location and surface characteristics completes the solution with minimum calculations. This is done at the pixel or voxel level in the display; therefore, no interpolation of shading is needed.

When the surface rendering is accomplished in this way using section planes that are relatively widely spaced, it often produces artefacts in the reconstruction that appear to be grooves parallel to the section direction (**Figure 58**). The use of relatively large voxels can produce rendered





Figure 57. Diagram showing the construction of a triangular tessellation on a surface formed by discrete height values for an array of pixels.

Figure 58. Rendered surface of a fractured pelvis created from sequential tomographic section images. Note the appearance of grooves and other surface artefacts due to the section plane spacing.



results that have an artificial blocky appearance (**Figure 59**). Applying image processing operations beforehand to range image data is often used to improve the resulting surface image. Smoothing with kernels that calculate a weighted average can produce Phong shading. Applying a median filter removes noise that would show up as local spikes or holes in the surface. The names of filters such as the rolling ball operator discussed in Chapter 3 come directly from their use on range images. This particular operator tests the difference between the minimum value in two neighborhood regions of different sizes and eliminates points that are too low. The analogy is that depressions which a ball of defined radius cannot touch as it rolls across the surface are filled in.

Rendering a surface (calculating its brightness according to its orientation relative to the viewer and light source) using the lookup table approach is fast, provided that the appropriate tables for different light source locations and surface characteristics have been calculated beforehand.



Figure 59. Rendered surface of a spine in which the coarse voxel size produces a blocky appearance.

The tables are not large and can be stored for a reasonable number of cases. The tables permit a nearly real-time animation of a rendered surface with a moving light source, which is another way to use the dimension of time to reveal 3D spatial information. Changing the surface characteristics, from diffuse to specular or from white reflection to colored, can be used to show multiband data.

Surface rendering of structures isolated within a voxel array can be enhanced by adding stereo views (**Figure 60**) and shadows (**Figure 61**). Rendered surface images have the appearance of real, physical objects and so communicate easily to the viewer, although, they obscure much of the information present in the original 3D image data set from which they have been extracted. More complex displays, which require real ray tracing, can make surfaces that are partially reflecting and partially transmitting so that the surface can be combined with volumetric information in the display. This presentation has somewhat the appearance of embedding the solid surface in a partially transparent medium, similar to fruit in Jello. Such displays can be dramatic in appearance and useful for communicating complex 3D information, but are too slow to generate and have too many variable parameters to be used for most explorations of complex data sets interactively.

Figure 60. CT slices of arteries in a human heart, combined and surface rendered, shown as stereo pair (courtesy of General Electric Co.).




Figure 61. Stereo view of rendered dendrites, with shadows.

Multiply connected surfaces

Rendering techniques are most needed for complex, multiply connected surfaces, since the topology of such surfaces cannot be studied in 2D images. Rendering these more complex surfaces is also possible but takes a little longer. **Figure 62** shows a series of 2D planes in a 3D data set from an ion microscope. The sample, a two-phase metal alloy, shows many regions in each image. It is only in the full 3D data set that the connection between all of the regions is evident. In fact, each of the two phase regions in this specimen is a single, continuous network intimately intertwined with the other. This cannot be seen even by using resectioning in various directions (**Figure 63**).

Volumetric displays of this data set can show some of the intricacy, especially when the live animation can be viewed as the rotation is carried out. **Figure 64** shows a few orientations of the data set using ray tracing to produce a volumetric display; viewed rapidly in sequence, these produce the visual effect of rotation. The complexity of this structure and the precedence in which features lie in front and in back of others, however, limits the usefulness of this approach. Isolating the boundary between the two phases allows a rendered surface to be constructed as shown in **Figure 65**. With the faceting shown, this display can be drawn quickly enough (a few seconds on a desktop computer) to be useful as an analysis tool. A complete smoothed rendering (**Figure 66**) takes longer (several minutes on the same computer) or requires access to graphics workstations.

Figure 62. Sequential images from an ion microscope, showing two phase structure in an Fe-45% Cr alloy aged 192 hours at 540°C (courtesy M. K. Miller, Oak Ridge National Laboratories, Oak Ridge, TN).





Figure 63. Section views of the 3D voxel array formed from the images in Figure 62: (a) stored brightness values along several arbitrary orthogonal planes; (b) stored values on a series of parallel planes with dark voxels made transparent.




Figure 64. Volumetric displays using ray tracing and absorption rules as the planes from Figure 62 are shifted to produce the effect of rotation. Viewing in rapid succession produces a movie.

The rendering of the surface follows the determination of the various surface facets. The simplest kind of facet is a triangle. In the example shown in **Figure 65**, a series of narrow rectangles or trapezoids are used to connect together points along each section outline. For features in which the sections are similar in shape and size, this faceting is fairly straightforward. When the shapes change considerably from section to section, the resulting facets offer a less realistic view of the surface shape.




Figure 65. Simple rendering of the boundary surface between the two phases in the specimen from **Figure 62**, by interpolating planar facets between the planes.





Figure 66. Two high-resolution rendered views of the boundary surface between the two phases in the specimen from **Figure 62**.

The greatest difficulty is dealing with splits and merges in the structure. This means that the number of outlines in one section is greater than in the next, and so the surface must somehow divide. **Figure 67** shows two ways to do this. In one, the junction actually lies in one of the planes. It may either be located manually or by various algorithms, such as dividing the feature normal to its moment axis at a point which gives area ratios equal to those of the two features in the next section. The rendered result is fairly easy to draw because no surface facets intersect.

The second method constructs surface facets from sets of points along the feature boundaries from the single feature in one section to both of the features in the next section plane. This technique moves the junction into the space between planes, and produces a fairly realistic picture, but requires more calculation. The intersecting planes must be drawn with a z-buffer, a computer graphics technique that records the depth (in the viewing direction) of each image point, and only draws points on one surface where they lie in front of the other.

The major drawback to this kind of surface rendering from serial section outlines is that the surfaces hide what is behind them, and even with rotation it may not be possible to see all parts of a complex structure. Combining surface rendering with transparency (so that selective features are shown as opaque and others as partially transparent, letting other structures behind become

Figure 67. Two ways to render the surface on serial section outlines where splits or merges occur: (a) dividing one plane into arbitrary regions which correspond to each branch; (b) continuous surfaces from each branch to the entire next outline, with intersection between the planes.




Figure 68. Combining surface rendering with partial transparency to show structures otherwise hidden by surfaces in front. The spheres which lie behind the frontmost surfaces are dimmed or shown in a different color. Notice that this is only partial transparency in that the front vertical column is completely opaque to the rear one, and the spheres are dimmed uniformly without regard to the actual amount of material in front of them.

visible) offers a partial solution. **Figure 68** shows an example in which one kind of feature (the spheres) has been given the property to show with a reduced brightness through any superimposed surface. This allows the spheres to be seen where they lie behind the columns. The dimming is not realistic (since the density of the columns is not known), however, and other structures behind the columns are completely hidden.

In the example shown previously in **Figure 65**, the boundary was determined by thresholding in each of the 2D image planes, since the plane spacing was not the same as the in-plane resolution. A certain amount of interpolation between the planes is required, which makes the curvature and roughness of surfaces different in the depth or z direction. For data sets with true cubic voxels, as for example the tomographic reconstructions shown in **Figure 23** and examples in Chapter 11, the resolution is the same in all directions, and greater fidelity can be achieved in the surface rendering. **Figure 69** shows the surface rendered image from that data set.

Multiple colors are needed to distinguish the many features present (about 200 roughly spherical particles) in this structure. This use of pseudo-color is particularly important to identify the

Figure 69. Rendered surface image of spherical particles from tomographic reconstruction (see Chapter 11, **Figure 51**) created by interpolating surface tiles between the slices shown in **Figure 23**, and assigning arbitrary colors to each feature.






Figure 70. Rendered surface image of a triple braided textile reconstructed from serial section images. Colors are assigned to each fiber to make it easier to follow them through the braid. A series of shades of each color are used to indicate surface orientation.

continuity of multiply connected surfaces. **Figure 70** shows another example in which the structural organization of the braid is much easier to see when each fiber is assigned a unique color palette.

The braided fibers in **Figure 70** emphasize that serial section reconstruction for 3D displays is certainly not restricted to microscopy and medical applications. The study of woven textiles and fibers used in composite materials uses the same methods (Gowayed et al., 1991). **Figure 71** uses color coding to identify the muscles and bones in a common everyday example of serial sectioning. The images were acquired by visiting the local supermarket and photographing each roast sliced from a side of beef. After aligning the images and thresholding them to delineate the various structures, the stack of slices can be rendered to reveal the 3D structure of the muscles and bones. Remember, though, that the volume, surface area and length of these structures can be determine much more efficiently by using stereological procedures to draw grids and count hits on the individual slices as discussed in Chapter 8.

Image processing in 3D

The emphasis so far has been on the display of 3D image data, with little mention of processing. Most of the same processing tools that were described in the preceding chapters for 2D images can




Figure 71. Serial sections through a side of beef:


- (a) one section with features numbered;
- (b) stack of slices;
- (c) rendered result, showing selected muscles with solid surfaces and several bones with dashed outlines.

be applied more or less directly to 3D images for the same purposes (Nikolaidis and Pitas, 2001). Arithmetic operations such as ratios in multiband data are used, for instance, in exactly the same way. Each voxel value is divided by the value of the voxel at the same location in the second image. This kind of operation does not depend on the images having cubic voxels.

Many processing operations that use neighborhoods (e.g., for kernel multiplication, template matching, rank operations, etc.), however, do expect cubic voxel arrays. In a few instances, a kernel may be adapted to non-cubic voxels by adjusting the weight values so that the different distance to the neighbors in the z direction is taken into account. This adjustment only works when the difference in z distance is small as compared to the x, y directions, for instance a factor of 2 or 3 as may be achieved in the confocal light microscope. It will not work well if the image planes are separated by ten (or more) times the magnitude of the in-plane resolution. And any departure from cubic voxel shape causes serious problems for ranking or template-matching operations.

In these cases, it is more common to perform the processing on the individual planes and then form a new 3D image set from the results. **Figure 72** shows a series of pseudo-rotation views of the MRI head images used above. Each slice image has been processed using a Frei and Chen operator to extract edges. These edges show the internal structure as well as the surface wrinkles in the brain. Furthermore, image processing was used to form a mask to delineate the brain (defined

Figure 72. Several views of the brain from the MRI data set. The skull has been eliminated, individual image planes have been processed with an edge operator, and the resulting values have been used with emission rules to create a volumetric display. Lateral shifting of the planes produces pseudo-rotation. Each pair of images can be viewed in stereo, or the entire sequence used as an animation. Structures are visible from the folds in the top of the brain to the spinal column at the bottom.



as the central bright feature in each slice) and isolate it from other portions of the image. The result is a series of slice images that show only the brain, processed to show the internal edges.

Another display trick has been used here. It is not clear just what volumetric display mode is appropriate for such processed images. Instead of the conventional absorption mode, in which transmitted light is passed through the data array, these images use the emission mode, in which each voxel emits light in proportion to its value. That value is the edgeness of the voxel as defined by the Frei and Chen operator. In other words, we see the edges glowing in space. In a live animation, or for those readers who can use pairs of the images to view the 3D data set stereoscopically, this image creates a fairly strong impression of the 3D structure of the brain. The same type of display of lines in space can be used to display contours within a 3D data set.

This illustration may serve as an indication of the flexibility with which display rules for 3D images can be bent. Nontraditional display modes, particularly for processed images, are often quite effective for showing structural relationships. There are no guidelines here, except for the need to simplify the image by eliminating extraneous detail to reveal the structure that is important. An experimental approach is encouraged.

The use of 2D processing of image planes in a 3D data set should be used with some care. It is only justified if the planes have no preferred orientation and are random with respect to the structure, or conversely if the planes have a very definite but known orientation that matches that of the structure. The latter situation applies to some situations involving coatings. When possible, 3D processing is preferred, even though it imposes a rather significant computing load. The size of neighborhoods increases as the cube of dimension. A kernel of modest size, say 7×7 , may be fast enough for practical use in 2D, requiring 49 multiplications and additions for every pixel. In 3D, the same $7 \times 7 \times 7$ kernel requires 343 multiplications and additions per voxel, and of course the number of total voxels has also increased dramatically so that processing takes much more time.

For complex neighborhood operations such as gradient or edge finding in which more than one kernel is used, the problem is increased further because the number of kernels must increase to deal with the higher dimensionality of the data. For instance, the 3D version of the Sobel gradient operator would use the square root of the sum of squares of derivatives in three directions. And since it takes two angles to define a direction in three dimensions, an image of gradient orientation would require two arrays, and it is not clear how it would be used.

The Frei and Chen operator (Frei and Chen, 1977), a very useful edge detector in 2D images introduced in Chapter 4, can be extended to three dimensions by adding to the size and number of the basis functions. For instance, the first basis function (which measures the gradient in one direction and corresponds to the presence of a boundary) becomes

$$\begin{array}{ccccccc}
 -\sqrt{3}/3 & -\sqrt{2}/2 & -\sqrt{3}/3 & & & & \\
 -\sqrt{2}/2 & -1 & -\sqrt{2}/2 & 0 & 0 & 0 & \\
 -\sqrt{3}/3 & -\sqrt{2}/2 & -\sqrt{3}/3 & 0 & 0 & 0 & +\sqrt{3}/3 \quad +\sqrt{2}/3 \quad +\sqrt{3}/3 \\
 & & & 0 & 0 & 0 & +\sqrt{2}/2 \quad +1 \quad +\sqrt{2}/2 \\
 & & & & & & +\sqrt{3}/3 \quad +\sqrt{2}/2 \quad +\sqrt{3}/3
 \end{array}$$

It should be noted that in three dimensions, it is possible to construct a set of basis functions to search for lines as well as surfaces. It remains to find good ways to display the boundaries which these operators find.

Three-dimensional processing can be used in many ways to enhance the visibility of structures. In **Figure 42**, the boron concentration was shown volumetrically using emission rules, although, the overlap between front and rear portions of the structure makes it difficult to see all of the details.

The surface rendering in **Figure 43** is even worse in this regard. **Figure 73** shows the same structures after 3D processing. Each voxel in the new image has a value that is proportional to the variance of voxels in a $3 \times 3 \times 3$ neighborhood in the original image set. These values are displayed volumetrically as a transmission image. In other words, the absorption of light coming through the 3D array is a measure of the presence of edges; uniform regions appear transparent. The visibility of internal surfaces in this "cellophane" display is much better than in the original, and the surfaces do not obscure information behind them, as they would with rendering.

The time requirements for neighborhood operations are even worse for ranking operations. The time required to rank a list of values in order increases not in proportion to the number of entries, as in the kernel multiplication case, but as $N \cdot \log(N)$. This assumes a maximally efficient sorting algorithm and means that ranking operations in really large neighborhoods take quite a long time.

For template-matching operations such as those used in implementing erosion, dilation, skeletonization, and so forth, the situation is worse still. The very efficient methods possible in 2D by using a lookup or fate table based on the pattern of neighbors will no longer work. In 2D, there are eight neighbors so a table with $2^8=256$ entries can cover all possibilities. In 3D, there are 26 adjacent neighbors and $2^{26} = 67$ million patterns. Consequently, either fewer neighbors can be considered in determining the result (e.g., just the six face-touching neighbors), or a different algorithm must be used.



Figure 73. Volumetric display of Boron concentration from SIMS image data set. The series of images shows pseudo-rotation by shifting of planes, using the local 3D variance in pixel values to locate edges. The magnitude of the variance is used as an effective density value to absorb light along rays through the voxel array.



All of the morphological operations (erosion, dilation, etc.) have direct generalizations in three dimensions (Gratin and Meyer, 1992). Normally, for practical reasons, cubic voxels are used. This means that consideration must be given to the difference between face-, edge- and corner-touching neighbors. In principle, the more symmetrical arrangements of voxels mentioned previously would allow simpler erosion/dilation rules and more uniform processing, but these are rarely used. Just as in two dimensions it is necessary to distinguish between 4-connectedness (touching pixels share an edge) and 8-connectedness (touching pixels share a corner), in three dimensions 6-, 18-, and 26-connectedness possibilities exist for voxels (sharing a face, edge, or corner, respectively). For practical reasons, many methods use 6-connectedness as the simplest definition.

This has particular impact on thresholding and feature enumeration. In many 3D data arrays, thresholding based on voxel value is a simple and effective method. But in complex structures methods based on region growing are widely used, in which the user identifies a seed pixel and then all connected voxels having values within some tolerance band are selected. Another approach to segmentation of features uses a 3D extension ("balloons") of the active contours ("snakes") used in two dimensions (Kaes et al., 1987; Cohen, 1991).

The 3D analog of a Euclidean distance map can be constructed by a direct extension of the 2D method, and has the same advantages both for improving isotropy and for distance measurement from surfaces or boundaries (Borgefors, 1996). Watershed segmentation in three dimensions, however, has only rarely proved satisfactory.

Skeletonization in 3D analogous to that in 2D would remove voxels from a binary image if they touched a background or "OFF" voxel, unless the touching "ON" voxels did not all touch each other (Borgefors et al., 1999; Halford and Preston, 1984; Lobregt et al., 1980). If touching is considered to include the corner-to-corner diagonal neighbors as well as edge-to-edge touching and face-to-face touching, then a minimum skeleton can be constructed; however, if a table for the 26 possible touching neighbors cannot be used, then it is necessary to actually count the touching voxels for each neighbor, which is much slower.

It should be noted that skeletonization in 3D is entirely different from performing a series of skeletonizations in the 2D image planes and combining or connecting them. In 3D, the skeleton becomes a series of linear links and branches that correctly depict the topology of the structure. If the operation is performed in 2D image planes, the skeletons in each plane form a series of sheetlike surfaces that twist through the 3D object and have a different topological relationship to the structure. The skeleton of linear links and branches in 3D has been formed by connecting the ultimate eroded points (UEPs) in each section. A different skeleton can be constructed by joining the branch points in the 2D sections.

Measurements on 3D images

As discussed in Chapters 8 and 9, it is clear that one of the reasons to collect and process images is to obtain quantitative data from them. This is true for 3D imaging as well as 2D, and consequently some additional comments about the kinds of measurements that can be performed, their practicality and the accuracy of the results appear appropriate.

Measurements are broadly classified into two categories: feature-specific and global or scene-based. The best known global measurement is the volume fraction of a selected phase or region. Assuming that the phase can be selected by thresholding (perhaps with processing as discussed in Chapter 3), then the volume fraction could be estimated simply by counting the voxels in the phase and dividing by the total number of voxels in the array or in some other separately defined reference volume. The result is independent of whether the voxels are cubic. In fact, the same

result can be obtained by counting pixels on image planes and does not depend in any way on the arrangement of the planes into a 3D array.

A second global parameter is the surface area per unit volume of a selected boundary. There are stereological rules for determining this value from measurements on 2D images, as presented in Chapter 8. One method counts the number of crossings that random lines (for a random structure, the scan lines can be used) make with the boundary. Another method measures the length of the boundary in the 2D image. Each of these values can be used to calculate the 3D surface area.

It might seem that directly measuring the area in the 3D data set would be a superior method that does not require so many assumptions. In practice, it is not clear that this is so. First, the resolution of the boundary, particularly if it is irregular and rough, depends critically on the size of pixels or voxels. The practical limitation on the number of voxels that can be dealt with in 3D arrays may force the individual voxels to be larger than desired. It was pointed out before that a 1024×1024 image in 2D requires 1 megabyte of storage, while the same storage space can hold only a $128 \times 128 \times 64$ 3D array.

The use of smaller pixels to better define the boundary is not the only advantage of performing measurements in 2D. The summation of boundary area in a 3D array must add up the areas of triangles defined by each set of three voxels along the boundary. A table can be constructed giving the area of the triangle in terms of the position differences of the voxels, but the summation process must be assured of finding all of the parts of the boundary. No unique path can be followed along a convoluted or multiply-connected surface that guarantees finding all of the parts.

For other global properties, such as the length of linear features or the curvature of boundaries, similar considerations apply. The power of unbiased 2D stereological tools for measuring global metric parameters is such that the efficiency and precision of measurement makes them preferred in most cases.

Feature-specific measurements include measures of size, shape, position, and density. Examples of size measures are volume, surface area, length (maximum dimension), and so forth. In three dimensions, these parameters can be determined by direct counting. The same difficulties for following a boundary in 3D mentioned previously still apply. But in 2D images the measurements of features must be converted to 3D sizes using relationships from geometric probability. These calculations are based on shape assumptions and are mathematically ill-conditioned. This means that a small error in measurements or assumptions is magnified in the calculated size distribution.

Simple shapes, such as spheres, produce good results. **Figure 74** shows the result for the tomographic image of spherical particles shown in **Figure 65**. The measurement on 2D plane slices gives circle areas that must be unfolded to get a distribution of spheres as discussed in Chapter 8. The result shows some small errors in the distribution, including negative counts for some sizes that are physically impossible, but the total number and mean size are in good agreement with the results from direct 3D measurement and require much less effort.

When feature shapes are more complicated or variable, 2D methods simply do not work. If information on the distribution of shapes and sizes is needed, then measurement in 3D, even with the problem of limited resolution, is the only available technique.

Position of features in 3D is not difficult to determine. Counting pixels and summing moments in three directions provides the location of the centroid and the orientation of the moment axes. Likewise, feature density can be calculated by straightforward summation. These properties can be determined accurately even if the voxels are not cubic and are affected only slightly by a reduction in voxel resolution.




Figure 74. Comparison of 2D and 3D measurement of size of spherical particles in structure shown in Figure 69:

- (a) size distribution of circles in 2D plane sections;
- (b) estimated size distribution of spheres by unfolding the circle data in figure a (note negative values);
- (c) directly measured size distribution of spheres from 3D voxel array.

Shape is a difficult concept even in two dimensions. The most common shape parameters are formally dimensionless ratios of size, such as $(\text{volume})^{1/3}/(\text{surface area})^{1/2}$ or length/breadth (length is easy to define as the longest dimension, but just as for the 2D case, the proper definition and measurement procedure for breadth is not so obvious). The selection of a parameter which has meaning in any particular situation is very ad hoc, either based on the researcher's intuition or on trial-and-error and statistical regression. In 3D, the values may be less precise because of the poorer voxel resolution, but the accuracy may be better because the size parameters used are less biased. It may also be important to the user's intuition to consider 3D shape factors, which are less unfamiliar than 2D ones.

The other approaches to shape in 2D are harmonic analysis (which unrolls the feature boundary and performs a Fourier analysis on the resulting plot), and fractal dimension determination; both were discussed in Chapter 9. These parameters can be determined rather efficiently in two dimensions, but only with great difficulty in three dimensions. The 2D results are related stereologically to 3D structure, so it is preferable to perform these measurements on the individual 2D image planes.

Closely related to shape is the idea of topology. This is a non-metric description of the basic geometrical properties of the object or structure. Topological properties include the numbers of loops, nodes and branches (Aigeltinger et al., 1972). The connectivity per unit volume of a network structure is a topological property that is directly related to such physical properties as permeability. It is not possible to determine topological properties of 3D structures from single 2D images (although minimal 3D probes such as the disector can provide unbiased estimates). These can be measured directly on the 3D data set, perhaps after skeletonization to simplify the structure (Russ and Russ, 1989). Another example, a reconstruction showing the topology of an overhand knot, was shown in Chapter 8.

Conclusion

There is little doubt that 3D imaging will continue to increase in capability and popularity. It offers direct visualization and measurement of complex structures and 3D relationships, which cannot be as satisfactorily studied using 2D imaging. Most of the kinds of imaging modalities that produce 3D images, especially tomographic reconstruction and optical sectioning, are well understood, although the hardware will benefit from further development (as will the computers and software). Current display methods are barely adequate to the task of communicating the richness of 3D image data sets to the user. New display algorithms and interface control devices will surely emerge,

driven not only by the field of image processing but also by other related fields such as visualization of supercomputer data and interactive computer games. The continued increase in computer power and memory is certain. Watching and using these developments offers an exciting prospect for the future.

Imaging Surfaces

In many different disciplines, surfaces are more important than bulk structures. Mechanical interaction between parts involves friction and wear between surfaces, many chemical interactions take place on surfaces (including catalysis) and most modern electronic devices consist of thin layers of materials laid down in intricate patterns on the surface of substrates. The appearance of objects is dominated by their surface characteristics, textures, and coatings. In all these cases and many more, scientists and engineers need to characterize surfaces and the ways in which fabrication and use modify them. Imaging plays important roles in obtaining the information as well as presenting it for human visualization and analysis.

Producing surfaces

Surfaces are produced in a wide variety of processes, some tightly controlled and some quite chaotic. One of the oldest techniques by which mankind has produced intentional surfaces is by removal of material, for instance creating a statue or a stone tool by removing chips from a larger block of stone. Modern fabrication of parts typically involves machining, grinding, and polishing to remove material and to create a surface with specific macroscopic dimensions and also microscopic roughness.

Machining is a process in which a cutting tool removes chips from the material as it is moved relative to the workpiece. The shape of the tool's cutting tip or edge, its speed and the depth of cut, control the dynamics of chip formation which can be either ductile (long, continuous chips) or brittle (short, broken ones). The surface typically displays long grooves in one direction whose shape is determined in large part by the shape of the tool. Grinding is a process in which many small cutting points, typically facets of hard particles cemented together into a wheel, simultaneously remove material from a surface. Polishing results when many loose hard particles slide and roll between two surfaces, removing material as the surfaces move relative to one another. Impact erosion (such as sand-blasting) uses particles to produce small craters on the surface. Each of these processes involve both plastic deformation and fracture, and have many variables such as applied forces, the presence of liquids, etc., which dramatically modify the appearance and performance of the resulting surface (as well as the tools or particles doing the work). A wide variety of other methods are used, ranging from fracture to electrical spark discharges, plastic deformation of surfaces by rolling, forging or extrusion, chemical etching, and so on, which modern technology employs to produce surfaces of parts by the removal or rearrangement of material. **Figure 1** shows a few different surfaces.




Figure 1. Range images of metal surfaces (each shows 1 square millimeter with grey scale proportional to elevation):

(a) machined (flycut) surface of aluminum;

(b) ground surface of stainless steel;

(c) vapor-polished surface of aluminum;

(d) shot-blasted surface of brass. Images courtesy Rank Taylor Hobson Ltd., obtained with a scanning probe instrument with 5 μm radius diamond stylus.

Other methods can build up surfaces by deposition. Again, this may be physical or chemical. Liquids solidify to leave solid coatings, sometimes accompanied by polymerization or formation of crystalline structures. Liquids may solidify in a mold which controls some of the surface morphology, while in some cases other forces such as viscosity and surface tension are more important than the actual mold surface. Freezing of liquids or gases onto a substrate may produce either very smooth or extremely rough surfaces depending on how the particles and molecules can move on the surface (**Figure 2**). Electroplating typically produces quite smooth surfaces by deposition of



Figure 2. Photograph of a glaze covering a ceramic pot. The glaze flows down the surface in a molten state and solidifies to an amorphous glass under the forces of surface tension. Subsequently, the atoms rearrange themselves to form crystals which nucleate and grow in the glaze, producing visually interesting patterns and also modifying the surface geometry.

atoms from chemical solution, while ballistic deposition and aggregation of particles may generate ones with porous fractal surfaces. Some deposited layers are then subjected to selective removal, either by chemical or physical processes. This is the process by which complex multilayer electronic chips and a coming proliferation of micromechanical devices are fabricated.

Another concern about surfaces is their cleanliness. The presence of particulates, either lying loosely on the surface or attached by electrostatic or chemical forces, can disrupt the deposition of the carefully controlled layers used in microelectronics, so elaborate clean rooms and handling methods are required. Surface defects such as pits and scratches are also of concern. Chemical modification of surfaces is called contamination, oxidation, or corrosion depending on the circumstances. This is strongly controlled by the environment. Sometimes such processes may be carried out intentionally to protect the original surface from other environmental effects (for instance, aluminum is anodized to produce a thin oxide layer that provides a chemically inert, mechanically hard surface resistant to further contamination in use). Electrical and optical properties of surfaces can be modified greatly by extremely thin contamination layers.

In all these cases, there is a great need to characterize the surfaces so that the topography of the surface, and perhaps other properties such as chemical composition or electrical parameters, can be determined. Some of the surface characterization data are obtained directly by imaging methods. Even when the data are obtained in other ways, visualizing the surfaces is an imaging technique, relying on the human interpretation of the images to detect important information about the surfaces. Measurement follows to reduce the image data to a few selected numbers that can be used for process control, and to correlate the structure of the surfaces with their fabrication history on the one hand and with their performance and behavior on the other.

Devices that image surfaces by physical contact

Most of the measurement methods used to characterize surfaces are based on either some kind of microscope that provides magnified images of the surface, or scattering of radiation or particles from the surface. The methods may provide measurement of either composition or geometry, including the thickness of thin layers. Many different kinds of microscopes (and some tools that may

not be conventionally thought of or named as microscopes) are used to study surfaces (Castle and Zhdan, 1997; Van Helleputte et al., 1995; Russell and Batchelor, 2001). Many of these require no surface preparation, or at most minimal cleaning; a few such as the scanning electron microscope (SEM) may require applying conductive coatings to electrical insulators. The common methods use visible light, electrons, ions, physical contact, electron tunneling, sound waves, and other signals to produce images that sometimes can be directly related to the surface geometry and in other cases are primarily influenced by the surface slope, composition, coating thickness, or microstructure. The most immediately useful imaging methods are those where the output consists of "range" values in which the elevation of the surface is directly represented in the values often shown as either profile traces or grey scale images. Although most of the examples shown here will be ones in which the grey scale directly encodes elevation, it should be understood that a similar display of chemical information or elemental concentration can be dealt with using the identical measurement and visualization tools.

One method that covers some of the newest devices, such as the atomic force microscope (AFM), and quite old and well-established methods used in industrial manufacturing, such as profilometers, is the use of a mechanical stylus that is dragged across the surface. Motion of the stylus is amplified to record the elevation of the surface point by point. If a full raster scan is used, this produces an array of elevation data that can be displayed as an image, as shown by the examples in **Figure 1**. If the mass of the moving parts of the stylus assembly is kept as low as possible, forces of a few milligrams can keep the stylus in contact with the surface (at least for surfaces that have slopes up to about 45°) at quite high scanning rates. The images in **Figure 1** were obtained in about 50 seconds each, as an array of 500 × 500 points covering a 1-mm square area.

Stylus instruments used in industry typically use diamond-tipped styli with a tip radius of about 1μ, which defines the lateral resolution that the instruments can provide. Vertical motion may be sensed using inductive, capacitance or interference gauges, which are capable of sub-nanometer sensitivity. With suitable calibration, which is typically provided by scanning over known artefacts, these methods are routinely used to quantitatively measure surface elevations in many industrial settings to measure surface finish, the thickness of layers, etc. These instruments have primarily been used with metal and ceramic parts, but are also capable of measuring a wide variety of softer and more fragile materials as shown in **Figure 3**.

The AFM also uses a stylus but a much smaller one (Quate, 1994; Binnig et al., 1986; Wickramasinghe, 1991). The scanning tunneling microscope (STM) stimulated a range of new microscopies which use essentially the same scanning and similar feedback principles to obtain nanometer resolution images. The AFM was introduced in 1986 as a new instrument for examining the surface of insulating crystals, and it was clear from the beginning that it would be capable of resolving single atoms, although unambiguous evidence for atomic resolution did not appear until 1993. The AFM has evolved into a flexible instrument that provides new insights in the fields of surface science, electrochemistry, biology and physics, and new adaptations of the technology continue.

By etching silicon or silicon nitride to a sharp point, or by depositing carbon in such a way that it grows into a long thin spike, a stylus can be fabricated with a tip radius of a few nanometers (**Figure 4**). This allows much greater lateral resolution than profilometer styli. But such tips are extremely fragile and easily deformed, so a variety of techniques have been devised to utilize them to probe a surface. Typically, the tip is used as a reference point and the surface is translated in the z (elevation) direction to contact it. The tip is attached to or a part of a cantilever arm whose deflection is monitored by deflection of a light beam on the rear face or sometimes by interference measurement, and vertical sensitivity below one nanometer is easily obtained. Either the sample or the stylus can be translated in an X, Y raster pattern to cover the entire surface to create a complete image. The translation is typically accomplished with piezoelectric devices, which limits the total range of motion.




Figure 3. Range images obtained with a scanned stylus instrument: (a) paper; (b) microelectronic chip.



Figure 4. Scanning electron microscope image of an ultrafine tip used for high-resolution atomic force microscopy (image courtesy Topometrix Corp.).

The traditional and still most common mode of operation places the tip in sliding contact with the surface. In order to reduce the lateral and shear forces on the stylus and the surface, the stylus may be rapidly raised and lowered ("tapping mode"), or the lateral forces may be measured by the twisting of the stylus to determine the elastic modulus of the surface material, or the friction between the stylus and surface. Additional modes can be used in which physical contact is not actually required. For example, the stylus can track the surface without touching it with somewhat lower resolution by using attractive Van der Waals forces. In addition, some systems use strategies such as heating the tip and measuring the heat loss when it is close to the surface, vibrating it and measuring a change in characteristic frequency when it is close to the surface, or using it as a guide for light photons that interact with the surface and detect its presence without contact. The electric or magnetic force gradient and distribution above the sample surface can be measured

using amplitude, phase or frequency shifts, while scanning capacitance microscopy measures carrier (dopant) concentration profiles on semiconductor surfaces. The original operational mode, which won a Nobel prize for its development, was scanning tunneling microscopy (STM), which measures the surface electronic states in semiconducting materials. The variety of operational modes of the scanned probe microscope seems nearly unlimited as manufacturers and users experiment with them, but many of these techniques are applicable only to a particular set of materials and surface types. The same technology has been used to modify surfaces, either by pushing individual atoms around or by writing patterns into masks used for lithographic manufacture of microelectronic and micromechanical devices.

One of the problems faced by AFMs is the difficulty in making quantitative dimensional measurements. Most designs use open-loop piezoelectric ceramic devices for scanning, which suffer from hysteresis and non-linearity. Software correction, no matter how elegant, can only go so far in correcting the resulting image distortions and measurement errors due to its inability to adapt to the topography of each individual sample. This particularly affects the use of the AFM in the metrology-intensive semiconductor industry. A few recent designs use a much more expensive approach which employs a separate measurement device in each axis to provide a closed-loop measurement of the piezo scanner's movement. Using either interference or capacitance gauges, these permit accurate measurements to be made on small structures such as microelectronic and micromechanical devices, magnetic storage devices, and structures such as the compact disk stamper shown in **Figure 5**.

The AFM is limited in the area that it can scan and the speed with which it can do so, and in the relief of the specimen which can be present without interfering with the cantilever arm. Special designs which attempt to alleviate one or more of these limitations are required for specific applications, as is true for all surface measurement approaches. But it is useful to have an overview of the general range of capabilities of the different techniques. **Figure 6** shows of graph (a Stedman diagram named after Margaret Stedman of the British National Physical Laboratories) that plots the range of lateral and vertical distances which can be accommodated by various surface measurement techniques. Notice that the minimum vertical resolution and the stylus instruments have a much larger range. Some of the other techniques plotted on the diagram will be discussed later in this chapter, but none of them offers a perfect combination of range and resolution in both vertical and lateral directions along with quantitative accuracy and an ability to deal with most kinds of surfaces.




Figure 5. AFM image of a defect on a CD stamper (image courtesy Topometrix Corp.). The presentation modes are discussed in the text.




Figure 6. Comparison of the typical range and resolution of several surface imaging technologies.

Noncontacting measurements

Stylus instruments are limited in speed by the need to move a probe with finite mass across the specimen one line at a time. They also touch the surface which raises concerns about specimen damage. Indeed, AFMs have been used to create surface topography as well as to image it, and industrial stylus instruments are often accused of leaving surface markings where they have been used on soft metal surfaces. For some surfaces, the best solution is to make a replica that can be scanned. Plastic replicas can preserve fine detail as shown in **Figure 7**, and remove concern about damage to the original specimen.

It would appear that using light as a probe would make it possible to overcome concerns about speed or damage. Unfortunately, it also raises others. The principle drawback is that the light does not interact with the same surface that the tip feels, so that the measured elevation does not agree with that from the stylus methods (which are accepted according to various international standards for surface measurement). In many materials the light waves penetrate to a small distance



Figure 7. Scanned stylus image of a plastic replica of human skin:
(a) grey-scale representation of the elevation data;
(b) photorealistic visualization of the surface.

beneath the surface as they are reflected, and this surface impedance depends upon the dielectric properties of the material (which may be modified near the surface by contamination or oxidation layers). Very fine scale structure can also produce speckle and interference effects that alter the returned light in ways that mimic quite different surface structures and give incorrect results. Also, the presence of contaminant or oxidation films or local surface tilt angles of more than a few degrees can reduce the amount of light scattered back to the detector so that some points on the surface are not measured at all.

Light can be used to probe surfaces in a wide variety of ways. One choice is to use a beam of light focused to a point that is then scanned over the sample in the same way that a stylus would be, while other instruments image light from the entire area of interest at once. Point probes may use confocal optics to detect the distance to the specimen (also called focus detection), which also requires vertical scanning of either the sample or the optics. This can be done for each point, which is slow because of the need to move a finite mass, or the scan can be performed over the entire area for each Z setting as in most confocal light microscopes. In this method, the light intensity at each location is stored for each Z setting and then the peak intensity value (which may be interpolated between settings) is used to determine the surface elevation at that location. It is also possible to construct an optical point probe that uses a lens with large chromatic aberrations, and to detect the wavelength of light that is most strongly reflected. Because the lens brings each wavelength to focus at a different working distance, this provides a measure of the surface elevation. Other optical techniques such as triangulation are very sensitive to the local surface slope and have relatively poor lateral resolution.

The method which provides the greatest resolution over the greatest range is interference, and this can be done with either a point probe or for the entire surface, and with either monochromatic (usually laser) light or with white light. The classic Michelson-Morley interferometer uses mirrors to send light along two pathways, and then recombines them to produce interference patterns that show fringes corresponding to differences in dimensions that can be much smaller than the wavelength of the light. When one leg of the interferometer reflects light from the sample surface, these fringes can be used to directly measure surface elevations. The lateral resolution is only as good as the light microscope used to collect the reflected light, or about 1 μm , but the depth resolution can be 1 nm or better. For samples that do not reflect light well or that have steep cliffs or deep pits, however, there may not be enough light reflected, or the spacing of the fringes may be too close together, to provide results.

When monochromatic light is used it is possible to interpolate the elevation of a point to about 1 nm, about 1/1000th the wavelength of light; however, the fringes must be far enough apart (the surface elevation must vary gradually) that it is possible to keep track of the changes in elevation because the interference pattern repeats with every multiple of the light wavelength. Many modern systems use more than one wavelength of light or white light instead. This produces constructive interference only at one focal depth, where the path lengths are equal and all of the wavelengths are in phase. When a surface is being imaged, this means that only points along one iso-elevation contour are bright. Varying the distance between the specimen and the optics allows scanning in Z to determine the elevation of points over the entire image. This takes longer than a simple monochromatic interference pattern and is limited in precision of elevation values to the performance of the scanning hardware but can handle surfaces with much more relief and steeper slopes.

Indirect interference techniques such as projection of grids to produce moiré patterns, also produce two-dimensional arrays of elevation data. Imagine light streaming through venetian blinds onto the floor of a room. If the floor is flat, the strips of light will be straight when viewed from above. If there are irregularities, they show up directly as deviations in the lines of light and shadow. When it is scaled down to the dimensions of interest on surfaces or, at least, to distances of a few

micrometers, which is the resolution of the light optics used to view the stripes, this same structured light method is easily used to measure surface geometry. Image processing can be used to detect the edges of the shadows, interpolating along each scan line to accuracies much better than the pixel spacing. Depending on the geometry this can produce vertical measurement accuracy similar to the lateral resolution of the optics (typically about 1 μm), but, of course, for only a few locations across the sample surface unless the line pattern is scanned. Toolmakers' microscopes and quality control examination of planed surfaces of lumber (among other applications) use the same basic method.

In a modern modification of the technique, mirrors or prisms are used to deflect a beam of laser light in patterns across the work piece to produce the same type of image. This has the advantage of being able to measure in various orientations and directions. Closely related to the idea of structured light is shadowing of surfaces with evaporated or deposited metal or carbon coatings, followed by measurement of the shadows cast by features and irregularities on the surface. If the image of the grid pattern in the incident light passes through another similar grid, it produces a moiré pattern, with dark lines that can be used to reveal the shape of the object. This technique is particularly useful for revealing local strains and deviations of surfaces from ideal geometric forms. Because it is fast and noncontacting, this method is often used in medical applications, ranging from orthopedic work on curvature of the spine to measuring the curvature of the lens of the eye before and after corrective surgery.

Microscopy of surfaces

Most forms of microscopy produce images in which intensity is related to the reflection of light (or some other signal) from the surface. This is only indirectly related to the surface geometry, and includes other information such as composition. Despite the difficulties in interpretation, this is still the most widespread procedure for surface examination.

The standard light microscope at moderately high magnification has a comparatively shallow depth of field. This creates many problems for examining surfaces. If the surface is not extremely flat and perpendicular to the optical axis (for example, a metallographically polished specimen), it cannot all be focused at the same time. Only low-magnification light microscopes can be used to examine rough surfaces (e.g., for fractography), and do not give much information about surface geometry. The pattern of light scattered by rough surfaces under diffuse lighting can be used to determine the roughness. It has been shown (Russ, 1994; Pentland, 1984) that a surface with fractal geometry will scatter diffuse light to produce a fractal pattern, and that there is a relationship between the fractal dimension of the surface and that of the light. This has also been reported for scanning electron microscope images of such surfaces. But measuring the overall roughness dimension of the surface is not the same thing as determining the actual coordinates of points on the surface.


On the other hand, the depth of field of the conventional light microscope is too great to measure the important dimensions in the vertical direction on rough surfaces. Paradoxically, the confocal light microscope has a much shallower depth of field (and more importantly rejects stray light from locations away from the plane and point of focus), which allows it to produce true range images from irregular surfaces. In the confocal microscope, the image is built up one point at a time (usually in a raster pattern). Each image corresponds only to points at a particular focal depth, but repeating this operation at many focal depths produces both an extended focus image in which the entire surface is imaged, and a range image in which the elevation at each point is recorded. The resolution in both vertical and lateral directions is much worse than the scanned probe microscopes, but quite useful for many surface measurement applications including metrology of some microelectronic devices.

Because of its very large depth of field, coupled with excellent resolution (typically <10 nm, much better than the light microscope) the SEM is often a tool of choice for the examination of rough surfaces. Furthermore, the appearance of the secondary electron image that is most often recorded from this instrument looks reasonably “familiar” to most observers, who therefore believe they can interpret the image to obtain geometric information. Unfortunately, this is not at all simple. **Figure 8** shows an SEM image of a surface, sintered tungsten carbide particles. This is a relatively simple surface composed of uniform composition particles with relatively flat facets. But there is no unique or simple relationship between elevation or slope and the local pixel brightness. For relatively smooth surfaces, “shape from shading” methods can convert changes in intensity to changes in slope and thus extract the geometry. The influence of fine-scale roughness, edges, surface contamination, or compositional variation prevents this from being a general purpose approach. Backscattered electron imaging is less sensitive to many of these effects and is used for some metrology applications, but only gives “real” geometric dimensions when comparison to standards is available or extensive modeling of the interactions between electron and sample is performed.

The great frustration in using the SEM to examine surfaces is that while the images look quite natural to human viewers, and seem to represent surface elevation in a familiar way, determining actual dimension values from them is nearly impossible except in very constrained cases. Metrology of integrated circuits is used to determine lateral dimensions, but even in these cases the definition of just what the relationship is between the physical contour of an edge and the voltage profile of the signal is far from certain (and highly dependent upon the voltage used, the material being imaged, the detector type and location, etc.). Metrology is used for quality control in which consistency rather than absolute accuracy is important, and there is no attempt to extract measurements in the Z direction from such images (indeed, even the visibility of points near the bottoms of grooves or contact holes is a problem).

Stereometric imaging in which two (or more) different views of the surface are combined to measure elevation is the same in principle as the generation of topographic contour maps from aerial photographs (Wong, 1980; Wang, 1990). This is not an easy technique to automate (Wrobel, 1991; Heipke, 1992; Barnard and Thompson, 1980; Zhou and Dorrer, 1994; Abbasi-Dezfouli and Freeman, 1994) however, and even with careful control of imaging conditions and measurement of

Figure 8. SEM image of the surface of sintered tungsten carbide.



angles, the vertical resolution is typically much worse than the lateral resolution of the individual images. This is because the tilt angle δ between the two views must usually be small ($7\text{--}10^\circ$ is typical) to prevent points being hidden in one of the two views, and the angle enters the calculation as $1/\sin(\delta)$. The precision of lateral dimensions is magnified by any uncertainty in angle and limits the precision of the final result.

Measurement of the elevation difference between individual points is usually quite straightforward using a human to locate the same points in the two images. Then the parallax or offset of the points gives the elevation by straightforward trigonometry; however, to generate an elevation map for an entire surface requires matching a great many points, and requires automation to be practical. The two methods used for this are area-based or feature-based. Area-based matching uses cross correlation (either in the spatial domain or the frequency domain) to find the location of an area in the second image that most closely matches each area in the first. Changes in the visibility or contrast of the area between the two images, the presence of specular reflections, or repetitive structures that produce multiple matches, produce problems for this approach.

Figure 9 shows a typical result (from Chapter 12) in which two stereo images have been matched by cross correlation, testing each point in the left-eye view with possible points in the right-eye view to find the best match. The horizontal displacement (disparity or parallax) where the match was found measures the elevation of the point. Mismatched or unmatched points are typically present, and are filled in using a median filter. The calculated elevation values can then be used for measurement or visualization.

Feature matching detects locations in each image which have some characteristic such as a maximum local value of variance. These points are then matched against the similar list of points in the other image. This is generally more successful, but may match only a few thousand locations in the two images out of perhaps a million pixels, so that the intervening locations can only be interpolated. In all cases, constraints such as preserving the order of points from left to right and knowing the direction of tilt so that searching for matches need only occur in a small fraction of the total image area are important aids to the practical implementation of the methods.

Because of its great depth of field, the scanning electron microscope is very often used to obtain stereo pair images of surfaces. The SEM is also used to generate x-ray maps of surface composition, discussed separately in the next section. Instead of using electrons, light, or other radiation to form an image of the surface, quite a lot of information is available from the scattering or diffraction patterns that are produced. X-ray patterns contain data about the structure of either crystalline




Figure 9. Surface reconstruction from stereo pair images:
 (a) left and right views of a microfossil superimposed as a color composite for viewing;
 (b) rendered surface reconstruction.

or amorphous layers. Electron scattering patterns contain information on the crystallographic arrangement and also on local strains in the material (which shows up as the displacement or broadening of lines in the pattern). Scattered patterns of electromagnetic radiation are in effect the Fourier transform of the elevation profiles of the surface and their measurement is therefore a direct method to study the surface elevation in a way that separates the information on the form or figure (the intended large scale geometry of the part), the waviness (medium scale departures from the figure), and the texture or roughness (the fine scale details on the surface). In many cases, the scattering of reflected light may also be directly related to the intended use of the surface, for instance if it is a high-precision mirror. Images such as diffraction patterns can be processed and measured using many of the same techniques as more conventional images, but are not discussed here in detail.

Surface composition imaging


The variety of techniques for probing surfaces is astoundingly broad. Some of the more common ones are the SEM using an x-ray detector, the ion microscope or microprobe using secondary ion mass spectrometry (SIMS), and Fourier transform infrared (FTIR) spectroscopy. There are many other tools used as well, particularly for measuring the thickness and composition of coatings. These include the backscattering of particles, which samples the composition and density of the sample at depths up to several micrometers beneath the surface. Also of rather specialized interest is acoustic microscopy, which is sensitive to debonding between the coating layer and substrate. Gigahertz acoustic waves have a wavelength similar to visible light, and can be used to image surface and near-surface structures that are difficult to detect with other signals. Surface waves are strongly reflected by cracks (even closed ones that cannot be seen otherwise) and bulk waves are similarly reflected by the surfaces of pores (although these subsurface waves only propagate at lower frequencies, with correspondingly poorer resolution). The speed of sound in the material can also be measured to determine the modulus of elasticity and other physical properties. Ellipsometry takes advantage of the fact that for many types of thin layer coatings, the plane of polarized light is rotated as it passes through the coating. Measurement of that rotation can provide highly precise coating thickness measurements, and when different wavelengths of light are used (or a spectrometer is used to scan an entire range of wavelengths) can also reveal details about the internal structure of the coating. This method is primarily used to measure relatively large spots, however, and not to produce images of the surface. Similarly, another spot-analysis analytical technique uses a laser-beam directed at a selected point on the surface with a light microscope to vaporize material from a pit (typically several μm across and deep) blasted from the surface so that the atomic and molecular fragments can be weighed in a mass spectrometer.

Several different types of ion microscopes are used. Many can produce elemental composition maps of the surface, or a series of such images at various depths in the material. An incident beam of ions knocks the uppermost layer of atoms loose from the specimen, either one point at a time (the ion microprobe) or over the entire surface at once. These atoms are ionized and are then accelerated into a mass spectrometer which separates them according to their mass/charge ratio, identifying specific elements and isotopes. A detector or detector array then produces an image. This typically represents the spatial distribution of one selected element at a time across the imaged area, with a lateral resolution of about $1\ \mu\text{m}$ (depending on the diameter of the incident beam in the case of the ion microprobe and the resolution of the ion optics in the case of the ion microscope) but with a depth resolution of one atomic layer. Rapid switching from one element to another as layers are removed produces complete data sets of the structure of the material.

Compositional mapping of surfaces is particularly important for examination of deposited coatings and the identification of contamination. The most common approach to this mapping uses a raster-scanned electron beam to generate characteristic x-rays from the atoms present, which are

then detected. The unique energy or wavelength of the x-rays identify the elements, and calculations based on the physics of x-ray generation can be used to determine their amounts. The lateral and depth resolution is limited to the order of $1\text{ }\mu\text{m}$ by the range of the electrons. Several different types of x-ray spectrometers are used; the diffractive or wavelength-dispersive type measures x-rays from one element at a time but with good trace element sensitivity while the more common type energy-dispersive type can measure all of the elements present at the same time, but with poorer detectability. These typically produce "dot map" images for several elements at once (**Figure 10**), which only approximately delineate the regions containing the elements and must be processed and combined as discussed in the following section. Other signals, such as Auger electrons, come from a smaller region near the point of entry of the focused electron beam and have better spatial and depth resolution, but because the signal to noise level is poor are not so good at detecting minor and trace elements.

Figure 10. X-ray maps from a mineral (mica) showing the location of several elements (as labeled on each image). Each dot corresponds to the location of the beam when one x-ray with the proper energy was detected, but some of these x-rays are the Bremsstrahlung background and are not characteristic x-rays of the element selected. This produces a finite background level in parts of the image where the specimen may not contain the element. Only in high-concentration areas are the boundaries of phase regions well delineated (images courtesy Pia Wahlberg, Danish Technological Institute).



Molecular identification of coatings and contamination can be made using infrared spectroscopy, in which various vibratory modes of the molecules are excited to produce characteristic spectral peaks. This method is most suitable to the analysis of organic materials and coatings, such as plastics. The spatial resolution of this approach is limited to several micrometers by the light optics used.

Processing of range images

Elevation data from surfaces produced by the various methods discussed above are typically recorded as 8- or 16-bit grey-scale images. Each pixel has a value that represents the physical elevation or composition of the corresponding surface location. Because most of the techniques described look vertically down upon the surface, the data are single valued and represent only the uppermost point for surfaces whose irregularity is so great that undercuts and bridges can occur. The SEM is an exception to this, as shown in **Figure 11**, which shows a complex polymer surface with undercuts and bridges which a range image cannot reveal. Integer data stored for each pixel in a range image can be converted to an elevation value in appropriate units (nm, μm , etc.) using scale data that are usually stored in the file header. Unfortunately, no standard format is used for this data. Not only does each manufacturer have its own format (which is not always generally readable or well documented), but some have more than one format that corresponds to different instruments.

Reading in these different file formats and storing the data in some standardized format such as TIFF files may require custom programming. Some image processing and display programs do have the ability to read arrays of data (i.e., images) in a wide variety of data formats provided that the user can specify (or deduce) the necessary format information. This typically includes at least the length of the header and perhaps where specific scaling or other information is stored within in, the dimensions of the array and whether the data are stored in rows or columns, and the data format (byte, integer, long integer, real, etc., and whether the byte order is Intel or Motorola — low word first or high word first).

Once the data are available, the kinds of processing that are required for surface images depend strongly on what kind of instrument was used. Some examples will serve to illustrate the possibilities:

Interference microscopes often have drop-out pixels where the local slope of the surface was too great (more than a few degrees from normal) to return enough light to the optics to permit




Figure 11. Scanning electron microscope of the complex surface of a polymer, with undercuts and bridges.

measurement. These points can be detected by filling the array beforehand with an illegal or impossible value that is replaced by real measurement data. Any pixel that retains the original value is a drop out point and must be filled in, and the most common way to do this is with a median or smoothing filter. A simple median would suffice for single points, but in many cases regions several pixels across or arranged as a line corresponding to some step or ridge on the surface may be dropped out. In this case a modification of the median approach is necessary. For example, for each dropout pixel, build an array of the neighborhood pixels that are not dropouts. If this is not an empty list, rank them into order and find the median, and place it into the missing location. After this has been done for all of the dropout pixels in the image, see if any remain. If so, repeat the process until every point has been filled in. This will fill in even large areas from their periphery and is sometimes preferable to linear or spline interpolation between pixels around the drop out, which produces blurred and smoothed edges.

A slight improvement in the basic iterated median can be achieved by using a weighted median in which pixels contribute to the decision on the value to be entered into the central pixel in proportion to how far away they are. For a 3×3 neighborhood, the corner pixels are $\sqrt{2}$ farther away than the ones that share an edge. Placing each corner pixel value into the list of values to be sorted twice and placing the edge pixel values into the list three times gives a ratio of 1.5 which is close to 1.414... and produces the desired weighted median when the list is sorted. For a 5-pixel-wide circular neighborhood (**Figure 12**) the pixel distances are 1, $\sqrt{2}$, 2, and $\sqrt{5}$. Weights of 3 for the edge neighbors, 2 for the corner neighbors, and 1 each for the more distant pixels approximate the relative contribution of the neighbors according to their distance. Note that this weighting increases the length of the list of values that must be sorted, which slows the operation significantly and is why larger integer weights which would more closely match the inverse distances are not used. For a single dropout pixel with all 20 neighbors in a 5-pixel-wide circle contributing, the sorting list grows from 20 values to 32, approximately doubling the sorting time. If the list of values is even in length rather than odd, the average of the two central values in the sorted list may be used.

Reducing or removing noise in range images uses the same methods as other images: either a median filter or some type of averaging filter such as the Gaussian. For surface images, however, they are best not applied in the same type of round neighborhood as shown in Chapter 3. Instead, the neighborhood is restricted to pixels that lie on the same portion of the surface, excluding points at a different elevation or on a surface with a different slope. This is an example of an adaptive neighborhood filter, similar to the procedures used to smooth geographic data (also a range map) called kriging. In the example shown (**Figure 13**), the neighborhood restriction provides superior noise reduction while preserving fine lines and corners.

Stylus instruments, whether macroscopic ones with diamond tips several micrometers in diameter or atomic force microscopes using Buckytubes to probe much smaller lateral dimensions, share some of the same image analysis problems and require anisotropic filtering. The scan rate along each line (the X-direction) is typically determined by the dynamics of the stylus itself — the mass of the moving tip and the applied force — which determines the maximum speed at which the tip

Figure 12. Weighted median filter duplicates pixel values 2 or 3 times for the closest pixels (red) approximately in inverse proportion to their distance from the center, as shown in blue.





Figure 13. Noise reduction using a restricted neighborhood:

- (a) original image showing noise;
- (b) conventional median filter (7-pixel diameter) applied;
- (c) mean filter in a 9-pixel diameter and a threshold to exclude values more than 40 grey levels different from the local surface;
- (d) median filter with the same neighborhood as c.



can move across the surface while remaining in contact with it. Too high a force will result in damage to the surface or the tip, but too low a restoring force will allow the tip to skip over holes or fly from rising slopes. Depending on whether the stylus moves some sensing element (an interferometer, capacitance or inductance gauge for example, or perhaps just a beam of light) or the surface is moved to null the position of the stylus (the typical AFM mode of operation) and the signal to the piezoelectric drivers is recorded, the output signal for the surface elevation is usually an electrical voltage. This must be amplified and then digitized, and in the process suitable filtering can be applied with a time constant appropriate for the scanning speed. This reduces the noise along each scan line, and eliminates the need for subsequent digital filtering to be required to reduce noise. Filtering may be used as discussed in the following paragraphs to separate the low-frequency signals related to surface form and waviness from the high-frequency roughness value, but that is part of the process of measurement rather than image enhancement.

The situation is quite different in the Y-direction (from one line to the next). A significant amount of time passes between sequential lines in the raster scan, allowing for changes in the mechanical and electronic components. Most systems scan in one direction only (to minimize hysteresis problems) and have a retrace scan during which the stylus is raised and not in contact with the sample. Repositioning the stylus to the exact same value is very difficult when the resolution of these methods in the vertical direction is of the order of 1 nm. The result is that subsequent scan lines tend to be offset from each other. Because the eye is sensitive to abrupt changes in brightness that extend over large distances, this produces images in which a visible horizontal stripe pattern may be seen. Some AFM manufacturers attempt to alleviate this problem by adjusting each line so that the average value is the same as that of the preceding line. This is rarely a good idea — it means for example that if there is a rising peak somewhere in the image area, the background around the peak will be depressed as the peak rises, producing false data and even an incorrect visual impression of the surface, as shown in **Figure 14**.




Figure 14. Enhanced visibility for surface scratches by processing:
 (a) original range image of letter on the surface of a coin;
 (b) adaptive histogram equalization applied to a;
 (c) photorealistic surface rendering of a;
 (d) rendering of height information from a with surface brightness values from b.

A better solution can be used instead of mean or average value for this line-to-line adjustment, although it requires more computation. The ideal solution would be to align the mode values of sequential lines of data. Under the assumption that the surface consists primarily of a background level with some roughness superimposed on it, plus major features of interest that rise or fall with respect to that plane, the mode is by definition the most probable surface elevation value. For a relatively small collection of data points (most area scans have only a few hundred data points along each line), the mode is not robustly determined. But for any distribution the median is closer to the mode than is the mean. Just as the median value is preferred over the mean for filtering noise from an array of pixels, so the median offers a workable solution for adjusting the scan lines in a raster scan stylus image. **Figure 15** shows this method applied to a typical image of a rough surface.




Figure 15. AFM scan of an etch pit in silicon, showing the line offsets in the raw image (a), the darkening artefact resulting from correction by adjusting the mean value of each line (b), and the improved result using the median value (c).

Another artefact in AFM images that is best avoided by proper attention to the hardware but which seems in principle to be correctable to some extent in software, is the deconvolution of tip shape from the images. As shown in Chapter 5, if the point spread function of an image can be measured, dividing the Fourier transform of the image by the transform of the point spread function can remove much of the smearing or loss of resolution, so that the inverse transform yields an improved image. Real AFM tips are far from perfect, exhibiting various departures from a ideal symmetrical point. Scanning the tip over a known artefact such as a circular disk (Jarausch et al, 1996; Keller, 1991; Keller and Franke, 1993; Markiewicz and Goh, 1994, 1995; Villarrubia, 1994, 1996, 1997) allows calculation of the tip shape, and permits this deconvolution. In practice, the tips have a short life and no two are identical, so that frequent recalibration is required; however, tip deconvolution is only important at the highest magnification and finest resolution levels.

Processing of composition maps

Most surface composition maps that are obtained by ion mass spectroscopy, x-ray energy spectroscopy, or other methods, suffer from low signal levels. As discussed in Chapter 3, this can in principle be rectified by collecting more data, but this is generally not desirable on economic grounds and sometimes is impossible because the analysis alters or consumes the surface. Thus, noise reduction methods, such as weighted smoothing or median filtering, are often applicable. **Figure 16** shows the Potassium x-ray map from **Figure 10** processed with a Gaussian smoothing operator, a conventional median and a hybrid median. The latter gives the best preservation of edges, corners and fine lines.

When multiple images are obtained of the same area showing the spatial distribution of different elements or other chemical data, it is very important to find ways to display them in combinations that will communicate the information to the observer, and to find ways to delineate and distinguish the various phases that are present. Combining multiple images as color planes offers one approach to this, as shown Chapter 1 and repeated in **Figure 17**; however, because of the way that the display hardware (and human vision) works, this allows only three planes (R, G, and B) to be assigned and there may be many more individual images available than that. The situation is analogous to the situation for remote sensing images; the Landsat thematic mapper records seven wavelength bands from the visible into the infrared, and other satellites capture even more. There is no straightforward way to "see" all this information at one time; the choice of which planes to show and in which colors can be quite subjective and can reveal (or conceal) quite different aspects of the information.

If the individual elemental maps can be thresholded to correspond to the intensity levels from individual phases, then Boolean combinations of the planes using "AND" and "NOT" permits forming binary images of each phase, which can then be measured. This process corresponds to setting up threshold ranges in an N-dimensional intensity space corresponding to the number of elements present, in which the ranges are rectangular prisms in shape. This is often adequate to distinguish




Figure 16. AFM scan of a chemically deposited surface before and after the correction of line offsets.




Figure 17. Reduction of random statistical fluctuations in x-ray intensity by processing:
(a) original Potassium image from Figure 10;
(b) Gaussian blur, $s = 1.5$ pixels;
(c) conventional median filter (5-pixel-wide circular neighborhood);
(d) hybrid median filter of the same size.

the phases present in real materials, but a more free-form shape that corresponded to the natural variations in intensity for each phase would be preferred.

Some statistical techniques plot the intensity of each pixel in each of the image planes as a vector or point in n -space (nine dimensions for the example shown here) and then search for clusters in that space. Once clusters are identified and the boundaries around them defined, the various phases present can be identified (MacQueen, 1968; Anderberg, 1973; Hartigan, 1975). This is a direct extension of methods discussed in Chapter 10. The pixels whose values lie within each cluster are then classified as belonging to the corresponding phase, and a new image can be generated with unique colors for each class so that the phases are delineated (Bright et al., 1988; Mott, 1995; DeMandolx and Davoust, 1997). **Figure 18** shows the results of an unsupervised cluster analysis technique using principal components analysis applied to the raw data from **Figure 10**; seven phases are identified and arbitrarily color coded.

When there is *a priori* information about the composition of the various phases expected to be present, this method works quite well. Cluster detection without such information, however, suffers from several problems. First, the statistical techniques will always be better able (in a statistical sense) to segment the space by defining more clusters, so unless the number of phase clusters is known the results are suspect. Second, clusters for minor phases representing only a few percent by volume of the structure will be represented by only a few percent of the points. Although these phases may be very important (e.g., for the properties of materials and the economics of mineral ores) they will be poorly defined in the n -space plot and very hard to detect. They are likely to be overlooked amid the background of points from pixels that straddle boundaries between major phases. Principal components methods, for example, are more likely to segment single phase regions based on minor gradients in composition or statistical variations in intensity, instead of identifying the presence of important minor phase regions.

Data presentation and visualization

Many types of surface measurement instruments produce data arrays with very large range to resolution ratios. In other words the number of bits that encode the elevation or other surface




Figure 18. Color combinations from the nine elemental x-ray maps in **Figure 10**. Each of the three planes (RGB) can be assigned to any of the nine planes, so this one data set can be displayed in more than 500 different ways.

characterization data are very large. Most image processing and display programs cope adequately with 256 grey levels, but even this exceeds the ability of human vision to distinguish them on a computer screen. Resolution of 1 nm over a range of 1 mm, which is quite possible with a high-precision stylus or interferometric instrument, produces a million levels (20 bits). This far exceeds the capabilities of displays or of perception. Consequently, the display routines must either select one part of the entire range to display, or compress the data to show the entire range. Processing can help, for instance by dealing with local slopes or derivatives rather than absolute values, but this also requires some user experience to interpret.

The very large dynamic range of the data from surface measuring tools can only be displayed on the computer screen at the cost of making small but measurably different value appear with indistinguishably small differences in brightness. One particularly effective way of assisting in the visualization of such small changes, often associated with dirt or defects, is the adaptive neighborhood histogram equalization introduced in Chapter 4. **Figure 19** shows an example, a range image of the raised surface of a letter on a coin. Scratches on both the lower coin face and the raised letter can barely be detected in the original, but become visually quite evident after applying the equalization procedure, which makes small differences larger while suppressing large ones. Other processing techniques, discussed in preceding chapters, can also of course be employed on range images, but should be done so always recognizing that the actual numeric elevation values are altered and can no longer be used for measurement purposes.

Range images, in which the grey scale value at each pixel represents the elevation (or some other surface parameter) at that point, contain all of the raw information in the data array. Even if the range can be accommodated by the display (for instance by dividing down the resolution with




Figure 19. Color-coding of the seven major phases in the mica sample from **Figure 10**, as identified by an unsupervised cluster analysis method.

which the data were acquired, using a local equalization technique, or creative use of false colors), the resulting image is an unfamiliar one to human observers, and requires experience to interpret. Using false colors to increase the ability to visually discriminate small changes makes the resulting images even more unfamiliar.

Contour maps draw isoelevation lines, which are exactly the same as topographic maps of the earth's surface. These are familiar to many people and can be more easily be interpreted because they make it easy to follow the contour lines to identify the shape of protrusions and valleys, and to identify points at the same elevation. Of course, they also eliminate a great deal of information (the elevation data for all of the other pixels on the surface), but this is part of the simplification that makes interpretation easier. **Figure 20** shows the elevation of a familiar coin displayed as a grey-scale range image, one that has been color-coded, and one reduced to a small number of contour lines (which have also been color coded to make it easier to distinguish their elevation values).

Contour maps are less successful at communicating visual information when the lateral scale of detail is finer, or when the surface is very anisotropic, as shown in **Figure 21**. In these cases the individual lines are close together and hard to distinguish, and the lines do not tie together different areas of the surface very well. Whenever contour lines become close together because of the presence of fine detail or steep slopes, it is necessary to reduce the number of contour lines to help clarify the map.

Figure 20. Range image of a coin (from a scanning stylus instrument, data courtesy of Paul Scott, Rank Taylor Hobson Ltd.), displayed as
(a) a grey-scale range image;
(b) a color-coded range image;
(c) a color-coded contour map with ten isoelevation lines.





Figure 21. Range images and contour maps (with five isoelevation lines) for an injection molded polymer surface (a,b) and a ground metal surface (c,d).



The visual impression of surface relief can be improved by processing. A directional derivative creates an “embossed” appearance with light and dark contrast along edges and gradients. As discussed in Chapter 5, a convolution kernel of weights of the form

$$\begin{array}{ccc} +1 & +1 & 0 \\ +1 & 0 & -1 \\ 0 & -1 & -1 \end{array}$$

produces the effect of lighting the surface from the upper left corner and produces the effect of shadows which the eye interprets as relief. The direction should always be from the top; if bright edges appear on the bottoms of edges the human vision system (which is accustomed to lighting from above) inverts the interpretation of the data and perceives hills as pits and vice versa.

The derivative image shows fine detail but hides the overall elevation changes in the data. This can be alleviated by combining the greyscale range image with the derivative. This may be done by adding the two (simply change the central value in the kernel from 0 to 1), but results that correspond more closely to the way vision perceives texture on surfaces can be obtained by multiplying the two images together. This is shown in **Figure 22**. A particularly attractive version of this display can be constructed by using the information from a color coded range image as well. In **Figure 22d** a hue, saturation, intensity (HSI) model was used with the elevation assigned to the hue for each pixel and the derivative assigned to the intensity (saturation is set to 50%). The shadows create an impression of relief while the color informs the eye about overall elevation values.

These results compare quite favorably to the results of a true rendering of the surface using each triangle of neighboring pixels as a facet and calculating the reflection of light from a light source in a fixed position as shown in **Figure 23**. In this type of calculation the surface can be given various reflectivity characteristics, either more diffuse or more specular. In the examples shown a full




Figure 22. Enhancement of the coin image from Figure 20:

- (a) directional derivative;*
- (b) the derivative added to the original grey values;*
- (c) the derivative multiplied by the original elevation values;*
- (d) color image with the elevation in the hue plane and the derivative in the intensity image.*

ray-tracing and Phong shading was used. The latter method varies the shade across each facet according to the angle variation between neighboring facets, and produces a very smooth and realistic rendered surface as used in computer-aided design (CAD) workstations. The rendered image can also be color-coded, by using the grey-scale rendering of reflectivity as the intensity channel and the elevation as the hue channel as discussed above. Examples of this are shown in **Figure 24**.

Altering the displayed lighting and shading of surfaces can be an extremely powerful visualization tool for surface examination, taking advantage of the abilities to interpret surface images that humans have evolved in response to real-world experiences. Specular enhancement of surface appearance has been demonstrated by (Malzbender et al., 2001). They use multiple light sources (40 in the examples shown) to illuminate a specimen, which enables the local surface orientation at each location to be determined by shape-from-shading, also called photometric stereo. This data set is then represented mathematically as a polynomial map of orientation and texture, so that the appearance can be computed with altered surface reflectivity characteristics and any selected




Figure 23. Rendered surfaces of the coin (from **Figure 20**) and the polymer (from **Figure 21**), treating the surface as though it were a diffuse scatterer (plaster of Paris) or a specular one (shiny plastic or metal). The grid pattern visible in the coin image is an artefact of the scanner used to obtain the images, which becomes more visually evident in this display mode.

illumination direction. **Figure 25** shows the enhanced appearance of a Cuneiform tablet with synthetic specular shading computed from the estimated surface orientation.

With this data set, it is also possible to interactively alter the illumination direction by moving the mouse over the image. **Figure 26** shows four images from this procedure, in which fine details on the surface can be studied as the lighting is altered. For instance, one interesting note is the fingerprint left in the wet clay by the scribe who prepared this tablet 4000 years ago in Sumeria, which can be discerned as a series of ridges near the upper left corner of the tablet. Such fine details are generally not observable without this enhancement technique.

Rendering and visualization

The views shown previously all look onto the surface from directly above. This normal view shows all of the data, but is not the most familiar to a human observer. The use of computer graphics to




Figure 24. Color coded rendered surfaces using a hue/intensity combination:

- (a) coin (from **Figure 20**);
- (b) skin (from **Figure 7**);
- (c) paper (from **Figure 3**);
- (d) shotblasted metal (from **Figure 1d**).

render the surface from an oblique point of view (which many programs allow to be selected or even interactively rotated) produces an image that makes it easier to visualize the surface morphology, and perhaps to detect the features of interest. The oldest and simplest approach to this is to plot elevation profiles along some of the horizontal rows of pixels, displacing each line vertically and laterally to create the effect of a surface, and skipping some rows so that the lines are adequately spaced apart. **Figure 27** shows two examples of this. In the first, the coin image from **Figure 20**, the lines are well spaced and the surface slopes gradual enough that there are only a few places where lines are hidden (and erased). The result is a fairly easy surface to interpret. In the second example, the polymer image from **Figure 21**, the presentation is harder to interpret because so many lines cross each other and the overall morphology is obscured. Also in these displays the width of each line profile is the same, so that there is no perspective applied to the view. In this type of isometric presentation the human familiarity with the rules of perspective causes this constant width of the data to be misinterpreted as giving the array a wider apparent dimension at the back than at the front.

Figure 25. Specular enhancement of the image of a cuneiform tablet by computing the rendered reflection from the surface based on orientation data calculated by shape-from-shading using multiple images with different light source locations (image courtesy Tom Malzbender, Hewlett Packard Labs, Palo Alto, CA). Bottom half of the image shows the original surface appearance.



With the continued advance in computer graphics capabilities, in the form of more processing power and displays with more grey levels and colors, much more realistic presentations can be generated. Adding perspective also makes the data seem more realistic, and adding cross lines that connect together points on success line profiles breaks the surface up into an array of square or rectangular tiles that improve the interpretability of the surface morphology by showing slopes in the second direction. Increasing the line density provides more information, but still must omit many lines and rows of pixels in order to avoid overwhelming the eye with too many disappearing lines. Coloring in the tiles according to the elevation of the points provides additional cues to depth. The results (shown in **Figure 28**) represented the state of the art for desktop computers only a few years ago, but this is now easily performed on the typical desktop machine.

The most visually realistic presentation uses actual surface rendering to control the brightness of each facet on the surface. Instead of square tiles that must be bent to fit the four corner points (which in general will not lie in a plane), triangular tiles are simpler to deal with. Three corner points define the triangle and the orientation of the facet with respect to the line of sight and a light source location permits calculation of the intensity to be assigned to the facet. This can be done with complete photorealism given the time and computing power, but there will in general be a very large number of facets to render and faster methods are sought. We can perform this operation by shading each facet according to the product of its absolute height and its derivative to produce a very quick and visually realistic rendering.

Figure 29 illustrates this method. In **Figure 29a**, the edges of the individual triangular facets are drawn in for clarity. But with this method it is practical to create a facet for every pixel and its immediate neighbors, so the full resolution of the data set can be displayed. The result is shown in **Figure 29b**. In computer graphics (as used in CAD programs for example) it is common to apply shading to facets so that they blend in with their neighbors (Gourard or Phong shading) and to not reveal lines where they meet. But while that method is important for the small number of large facets encountered in CAD renderings, it is unnecessary for the tiny facets that cor-

respond to each pixel, because the facets may cover only 1 or 2 pixels on the display. This also speeds up the process.

It is useful to compare this method against the slightly simpler display procedure of using the elevation to shade the facets, or using false colors to indicate elevation, as shown in **Figure 30**. These methods produce much less realistic results for visualization, and require a more educated eye on the part of the user. It is important to understand that human vision is an important tool for examining surface images, since the presence of defects or other features of interest is usually far more readily discerned by an experienced observer than is possible with computer pattern recognition programs.

It is also possible to introduce color to these displays using the same procedure as shown in **Figure 24**, applying the elevation as a color in the hue channel and the slope or Phong-rendered values in the intensity channel, while drawing each facet in its appropriate place on the screen to generate a perspective-corrected visual representation of the surface geometry. This is far more visually interpretable than simply applying false color, as shown in **Figure 31**. Generating such displays requires only seconds on a typical modern desktop computer.

Using this type of presentation communicates a much more effective representation of the surface geometry to most users, even ones with some experience, than does the simple grey-scale range image. **Figure 32** shows this mode of presentation for the same surface images presented in **Figure 1**. Even though the range images contain all of the data, and the perspective-corrected visualizations actually obscure some of it, comparison suggests that the latter are more realistic in appearance and hence more useful for visual recognition of characteristics or defects.

Modern computer graphics is also capable of rapidly redrawing surface views from different viewpoints. This can be used in several ways. Generating two realistic renderings of a surface from slightly different points of view allows using human stereo vision to interpret the depth of a surface. Creating a series of such images from different viewpoints can be used to display a "movie" showing a fly-over across the surface, and with enough computer horsepower this can be done in real time as an operator manipulates a joystick to interactively control the flight path. Combining this with the stereo display (**Figure 33**) creates a virtual reality world in which the surface can be viewed in detail.

Analysis of surface data

There is no doubt that human visual examination of well-presented visualizations of the surface geometrical and compositional information offers a powerful tool for recognizing defects and other specific characteristics of the surface. For many purposes, numerical measures of the surface are needed, which can be used for control purposes and to correlate the surface geometry or compositional variations with the creation and processing history of the surface, and with its performance behavior. For these purposes, analytical methods relying on computer processing of the data are needed, and it is far from clear just what should be measured to provide effective parameters for any given requirement.

The traditional measures of surfaces include the thickness of coatings, and the geometric features of the surface geometry. Most of these values, although quite precise, are highly dependent on other factors such as the material composition (and spatial variations of composition), the particular measurement procedure used, and the size of the measured area. Most thickness measuring procedures and some elevation measuring instruments naturally average over a lateral distance that is at least several micrometers and often much more. This is much larger than the vertical resolution most techniques are capable of, and may hide important details of the coating. Sampling strategies must be employed to determine spatial uniformity.

Figure 26. A few images of the specularly enhanced surface of the same tablet shown in **Figure 25** as the location of the computed light source is interactively shifted. This happens in real time so that the observer can visually interpret the fine details (images courtesy Tom Malzbender, Hewlett Packard Labs, Palo Alto, CA).



a



b



c



d




Figure 27. Isometric line profile displays of the coin image from **Figure 20** and the polymer image from **Figure 21**.




Figure 28. Grid or mesh displays of the coin (**Figure 20**), polymer (**Figure 21**) and deposited surface (**Figure 16**).




Figure 29. Deposited surface rendered with triangular facets where brightness is the product of the elevation and slope:
 (a) large facets that are 6 pixels wide;
 (b) facets connecting each pixel and its immediate neighbors.




Figure 30. Data from the polymer surface:
 (a) elevation value used to shade the facets;
 (b) product of elevation and slope.




Figure 31. Presentation of elevation data from the coin image:
 (a) each facet shaded according to the product of elevation and slope;
 (b) each facet assigned a color representing elevation;
 (c) Using the hue to represent elevation and the intensity to represent slope;
 (d) using hue for elevation and the Phong-rendered surface values for intensity.

Dimensional measurement of surfaces is one application where coordinate measuring machines, stylus instruments and optical interference techniques are all used. In many cases very exact dimensions are specified in the design of the part, and so the measurement does not require an area scan or an image, but simply the proper alignment of the measuring tool with the component. However, as dimensions become small, as in the case of microelectronic devices, it may be necessary to acquire an image to locate the point where measurement is to be performed. **Figure 34** shows an elevation profile taken from one scan line of an AFM, from which highly precise measurements can be taken on the width and height of the steps present.

Because the AFM is a relatively slow device which has difficulty handling large parts, and because there is always concern about surface damage when a physical contact is made, many of the metrology measurements on these devices are presently made by SEM. (The light microscope was used with earlier generations of devices, but the dimensions are now too small for the wavelength of light to resolve them.) **Figure 35** shows an SEM image of two lines of photoresist on a silicon wafer. Unfortunately, the SEM image is only indirectly related to the surface geometry.

Figure 32. The surface data from the four surfaces in **Figure 1** and the two surfaces in **Figure 3** rendered to show perspective-corrected visualizations with facets shaded according to the product of slope and elevation. The graphics have been expanded in the vertical (Z) direction to increase the perception of roughness and relief in the data.




Figure 33. The surface rendering of **Figure 32f** performed from two slightly different viewpoints, which are combined to form a stereo pair.





Figure 34. AFM image of a lithographic test pattern used to select the location for a single line scan used for dimensional measurement. When used for this purpose, the AFM requires quantitative position sensing such as an interferometer, instead of relying on measuring the signals sent to the piezoelectric positioners.







Figure 35. Portion of an SEM image of two parallel lines of photoresist on a silicon wafer and the signal profile across them. The relationship between the physical profile and the signal depends upon the slope and roughness of the sides of the lines, the composition of the material, the electron accelerating voltage used, and the electron detector used.




Consequently the signal profile along a scan line (actually averaged over multiple scan lines to improve the signal to noise ratio) is difficult to interpret to determine the line width. The pitch or spacing of the lines can be determined with fair accuracy under the assumption that the lines have the same shape, and thus generate the same signal profile. Consequently selecting any reproducible characteristic of the signal — the peak, or the maximum slope, etc., can be used to measure the distance between the lines. In order to measure the peak width accurately, however, there must be some absolute determination of where the edge lies (and even what that means given the slightly irregular shape of typical lines).

Computer modeling of the process of generating the SEM image signal can be carried out for various specimen geometries (and as a function of composition, electron beam voltage, and detector characteristics and placement). This is a time consuming process but still easier than fabricating physical standards for comparison. Even so, little accurate metrology is done in reality. Most manufacturers that use SEM images for metrology select some arbitrary feature of the signal that can be easily and reproducibly measured, such as the point of maximum slope or halfway between the darkest and lightest signal levels, and use that to monitor changes in dimension but without trying to determine the actual dimension. This is the classic difference between accuracy and precision, and works adequately for production control but is not adequate to support the development of new geometries and devices.

Profile measurements

Unlike the SEM, most instruments considered here do produce actual physical elevation profiles. Surface measurements have historically been assessed from these elevation profiles rather than using full two-dimensional images (because the instrumentation is simpler and the time required is much less, and hence because familiarity with the methods became established). By applying filters to the data (either digitally or in the amplifier electronics), different ranges of frequencies in the profiles can be separated which are traditionally described as the figure or form, waviness, and texture or roughness (**Figure 36**). Form is the overall gross geometrical shape, which is generally specified in engineering drawings, controlled by set dimensions, and described by conventional Euclidean geometry. The medium frequencies are called waviness and the high frequencies the texture or roughness. In machining processes, waviness is assumed to result from vibrations or deflections in the machine, while roughness results from more local interactions between the tool and the local microstructure in the material. These divisions are somewhat arbitrary and may differ

Figure 36. Deviations from intended shape are commonly divided into roughness (fine scale), waviness (intermediate scale) and form (large scale).



according to the size of the part. The cutoff frequencies used to define the filters are typically set to wavelengths from about 0.25 mm up to several mm to separate waviness from roughness. International standards specify these as part of the measurement procedure for many mechanical engineering applications.

Filtering to separate roughness from waviness and form data was originally done using analog RC filters in the electronics. In modern systems digital processing is used, with a least-squares line or arc fitted to remove the form and spatial Gaussian filter to separate the waviness and roughness. This generalizes directly to area scans which can also be filtered with an equivalent Gaussian filter, or the Fourier transform of the image can be filtered to select the desired range of frequencies. The form data are most often separated by least squares fitting of a plane, or some other Euclidean shape such as a cylinder or sphere that corresponds to the known intended form, or a generalized polynomial. **Figure 37** shows a simple example of form removal, in which roughness on a




Figure 37. Range images and surface visualizations of a ball bearing surface showing the overall form (a,c) and the results of flattening the data by subtracting the spherical shape (b,d). Grey-scale and vertical expansion of the data is possible after form removal.

ball bearing (spherical) surface is made more evident visually and also becomes easier to measure after the general curvature is subtracted. Deviations are then measured from the nominal form of the ball.

The roughness of surfaces is typically determined from the roughness profile after the low(er)-frequency components have been removed. A wide variety of measurement parameters are used, some of them codified in various International Standards Organization (ISO) or other international standards, and some of them corresponding to specific industries or equipment manufacturers (Rosen and Crafoord, 1997). A complete review of instrumentation and methods is in Whitehouse (1994), and up-to-date reviews of analysis procedures are covered in Thomas (1999) and Mainsah et al. (2001). The most widely used procedures perform statistical analysis on the elevation data without regard to its spatial arrangement. Examples include the maximum peak-to-valley range of elevations along the profile, the average absolute value of the deviation from the mean (Ra), or the statistical standard deviation of the elevation data (Rq). Another measure of the magnitude of the roughness is the difference in elevation between the five highest peaks and five lowest valleys (Rz), but this requires defining a peak or a valley. This problem becomes more difficult when applied to area scans or images.

Information on the spatial distribution of the elevation data includes parameters such as the number of peaks along the profile and the correlation length. The latter may be defined as the average distance between successive peaks, or between points at some specific elevation such as the mean elevation line left by removing the form and waviness. A more general definition of the correlation length comes from a plots as shown in **Figure 38**; this is just the magnitude of the autocorrelation function which can be determined from the Fourier transform of the profile. The autocorrelation function (ACF) is also of interest because, for surfaces produced by a large number of independent events (shot blasting, grinding, ballistic deposition, etc.), it has the same shape as the ACF of the "average event" that produced the surface.

Functional parameters are also used, which are presumed to correspond to particular usage of the surfaces. The "Abbott curve" is simply the cumulative histogram of the elevation data (**Figure 39**); it gives the area of contact which would be obtained by removal of a portion of the surface, either by in-service wear or by an additional fabrication step such as plateau honing of automotive cylinder liners.

Another approach called "motif," originally introduced in the French automobile industry and now used throughout Europe (Dietzsch et al., 1997), simplifies the profile to just the peaks which would contact another surface based on the height of the peaks relative to the intervening valleys and the width of the valleys. **Figure 40** shows the principle. Peaks are characterized by their depth (the height above the valley) and their separation distance. Peak and valley patterns are then combined according to their separation distance and depth to eliminate the small peaks on the sides of larger ones, until a minimum is reached which contains just the most important peaks.




Figure 38. The correlation plot shows the probability that points will have the same elevation value as a function of their lateral separation. The correlation length is defined as the point at which this plot drops to $1/e$ or 36.79%




Figure 39. The histogram (black) of the shotblasted surface range image (**Figure 1d**) and the same data shown as a cumulative histogram (red). This plots the fraction of points on the surface whose elevation is less than the value on the horizontal axes, and is called the Abbott–Firestone curve.



Figure 40. The basics of Motif combination — each peak and valley motif is measured by the depth (height of the smaller peak above the valley) and width. Motifs which have peaks smaller than the neighbor on either side, a width less than an arbitrary cutoff, and a depth less than 60% of the largest depth in the profile are combined with their neighbors to reduce the number of motifs present.

All these methods have serious limitations. They are highly dependent on the length of the profile scanned and the lateral resolution of the data points. They involve some very arbitrary definitions of what constitutes a peak or valley, ignore the fact that the profile path will not cross the highest or lowest points of most surface peaks and valleys, and do not correlate very well with the real subjects of interest, which are the processes by which surfaces are produced and their behavior for whatever service they are used. They are primarily suitable for specific quality control applications in which the real meaning of the parameters is hidden but consistency of measurement results can be used to keep a working process in control. Furthermore, they are very difficult to generalize to surface images produced by area scans of elevation.

All profile methods suffer from the fact that most real surfaces are not isotropic but have some directionality that results either from the way the surface was generated, the characteristics of the material itself, or the use it has been subjected to. This so-called “lay” of the surface can be simple (e.g. the ground surface in **Figure 1** is highly directional) or very complex and subtle. Measuring a profile perpendicular to the principal lay direction is the recommended approach, but for complicated surfaces this misses much of the actual character of the surface.

Because the history of profile measurements had generated (or accumulated) a rash of parameters, an effort is being made to rationalize the measurement of area scans. Supported by the ISO committee and spearheaded by researchers at the University of Birmingham (Stout et al., 1993), a set of statistical, spatial and functional parameters have been proposed which will probably evolve to form the basis of future international standards. These still contain some of the same limitations as the profile measures, such as the need to define what constitutes a peak and a strong dependence upon the size of the scan area and the lateral resolution of the points. Also, they do not include some of the potentially important methods such as topographic analysis, envelope or motif analysis, and fractal geometry. Because they represent an important starting point for surface description, some consideration of them is appropriate.

The Birmingham measurement suite

Four classes of measurement parameters are proposed, ones that deal with the elevation values without regard to their location (called amplitude parameters), ones that deal with lateral distances on the surface (called spatial parameters), ones that combine these together (called hybrid parameters), and ones that are believed to have some direct correlation with surface history and properties (called functional parameters). Within each group only a very few parameters, ones which have the most direct relationship to the more widely accepted profile measurement parameters, are selected. The symbols proposed for these parameters use the same nomenclature as those for profiles except that S (surface) is substituted for the R used in profile standards.

The amplitude parameters are simple extensions to area scans of the statistical measures that are used with profile plots. For instance, Sa is the analog to Ra , the arithmetic mean deviation. For an area scan it is the arithmetic mean of the absolute values of the elevation values from the mean plane (fit as discussed previously). Sa is preserved only because Ra is widely used, and that is so in turn because it was comparatively easy in the pre-computer days to design instruments to measure it. The root-mean-square deviation of the elevation points is a more robust measure, which is simply the standard deviation of the distribution of the elevation values, Sq . The variance (the square of the standard deviation) is the second moment of the distribution. The third and fourth moments are the skew and kurtosis, respectively, and these are also used as amplitude measurement parameters, called Ssk and Sku , respectively. For simple distributions which are not bimodal, these three parameters offer a reasonably compact statistical description of the surface heights.

The histogram of the surface elevation data (examples are shown in **Figure 41**) shows the overall range of surface elevation. The skew in the distribution distinguishes such cases as the narrow and deep grooves that may be important for distributing lubricant on plateau-honed cylinder liners in automobile engines (**Figure 41c**). In this case most of the surface has a very narrow range of elevations but the grooves, which cover only a small fraction of the area, reach down to much





Figure 41. Histograms of elevation values for a few metal surfaces:

- (a) shotblasted brass, which has a symmetrical distribution;
- (b) ground stainless steel, which has a slight negative skew due to the presence of a few deep but separated parallel grooves;
- (c) a plateau-honed cylinder liner with a negative skew resulting from the deep intersecting grooves which distribute lubricant.

Figure 42. The 20 highest peaks (red) and lowest valleys (green) on the shotblasted (a), ground (b), and plateau-boned (c) surfaces shown in Figure 41.



lower depths. A skew in the opposite direction would correspond to a surface with just a few high peaks or ridges rising up from a relatively smooth surface. But the histogram by itself contains no information on the spatial arrangement of the pits and valleys or the peaks and ridges. The same histogram would result from a surface with all of the high points collected together in one continuous ridge or distributed as thousands of tiny separate peaks. The properties of these two extreme surfaces would be quite different.

Just as for profiles, these statistical measures of amplitude are sensitive to the size of the sampled area. For most surfaces the standard deviation increases with the number of points measured, in fact for a fractal rough surface the slope of a curve plotting the variance as a function of size on log-log axes is one of the ways used to measure the fractal dimension.


For profiles, the parameter R_z is the difference in elevation between the average of the five highest peaks and five lowest valleys. For an area scan of a surface this is generalized to S_z , the difference between the ten highest peaks and ten lowest valleys. This is not purely an amplitude parameter, however, because it depends critically on the definitions of a peak and a valley. They cannot be simply the highest and lowest points on the surface (or pixels in the surface image), because these could be (and often will be) adjacent to each other and would all represent a single peak and valley. For a profile, the presence of a low point separating two high points might be taken to indicate separate peaks. This is a flawed definition because infinitesimal irregularities

should not be considered significant, and so some criterion for the depth of the valley between the peaks is required; but, on an area scan of a surface, even more is needed because the peak (or valley) covers an area and two or more local peaks may connect along intricate paths (a ridge) to be considered part of the same peak (and vice versa for valleys). In tracing this connectivity, it matters whether pixels are considered to touch all eight of their immediate neighbors or only the four that share edges with them.

There is much more information in the identification of peaks and valleys than just the S_z parameter that is the elevation difference between the average of the ten highest and ten lowest. Different surfaces give rise to very different shapes for peaks and valleys, and their sizes and shapes, orientation and spacing may all contain important characterization information. In **Figure 42** several examples are shown in which peaks are defined as eight-connected (pixels touch eight neighbors), and are required to be distinct down to 80% of the height of the peak. Valleys are defined in the same way. In this example the 20 highest peaks and lowest valleys are found. The method is similar to the “flood fill” algorithm used in image processing, starting with the highest local maximum (and proceeding down) and including all touching pixels that extend down to the 80% limit, while checking to see if the peak merges into an existing labeled peak. Notice that for the shot-blasted surface the valleys are relative smooth in outline while the peaks are very irregular. Also, for the ground surface the peaks (ridges) tend to be broader than the valleys (crevices), and for the honed surface the peaks are very large while the valleys are much smaller. All of these differences are consistent with our understanding of how such surfaces are produced, and they may give important insights into other surfaces and their functional performance.

Another parameter involving the peaks present in the surface is the number of them per unit area, called S_d s. Again, this depends upon the definition of a peak as discussed above. It is likely that secondary information about the peaks will also be important in a variety of applications. For instance, the uniformity of spacing of the peaks may play a role in cases where the surfaces are involved in electrical or thermal contact, friction and wear, or to judge the visual and aesthetic appearance. As discussed under image measurement, the mean nearest neighbor distance can be used to determine the tendency toward uniform spacing or clustering by comparing the value to the mean distance that a Poisson random distribution of the same number of points per unit area would have. **Figure 43** shows the surface of an injection molded polymer in which the peaks are relative evenly spaced (a complex function of the surface finish of the die, the temperature, pressure and viscosity of the polymer, and its molecular weight). This uniformity coupled with a spacing between peaks that is close to the spatial resolution limit of human vision produces an aesthetically pleasing appearance for the product.

Figure 43. The highest 20 peaks and lowest 20 valleys in a one square millimeter area of the molded polymer surface from **Figure 21**. The peaks are color coded from purple to red according to height and the valleys from green to cyan. Notice that one valley completely surrounds a small peak, and that the peaks are much more regular in shape than the valleys.




Note in the figure that the valleys have a very different shape and distribution than do the peaks. In some other applications the same information about the density and uniformity of pits rather than peaks would be of interest. An example is the surfaces of plates used to retain ink for printing applications.

Surfaces with anisotropy or lay can be characterized by spatial parameters derived from the autocorrelation function. As discussed in the chapter on Fourier space processing, the autocorrelation function (ACF) is obtained by squaring the magnitude of the complex variables in the Fourier transform while setting the phase to zero, which removes all spatial location information. The inverse transform produces the two-dimensional spatial image of the ACF. The parameters defined from this function are the texture aspect ratio Str , the texture direction Std , and the autocorrelation length Sal . Understanding these may be helped by examining the ACF images in **Figure 38**. The autocorrelation length is defined in the Birmingham report as the shortest distance in which the magnitude of the ACF drops to 20%. For the examples shown, this is the minimum radius of the contour line drawn at the 20% intensity level. The texture aspect ratio is the ratio of the minimum radius to the maximum radius, and the texture direction is the orientation of the maximum radius.

For the examples in **Figure 44**, the ACF of the polymer surface (**Figure 21**) is quite isotropic (indicating that the surface is also isotropic), so the aspect ratio is unity and there is no pronounced direction. The ground surface (**Figure 1b**) on the other hand has a strong preferred orientation

Figure 44. Autocorrelation function calculated for: (a) the polymer surface from **Figure 21; (c) the ground surface from **Figure 1b**; (e) the flycut surface from **Figure 1a**, with superimposed contour lines (b, d, f) indicating the shape of the function and the distance at which it drops to 20%.**



which is evident in the ACF and can be measured there. For the flycut surface (**Figure 1a**) the texture is more complicated, as is indeed evident in the original image which shows two predominant machining directions. The *Str*, *Std*, and *Sal* values as defined can of course be measured from the ACF but it is not clear that they contain all of the information about the surface lay that would be desired for characterization.

The hybrid properties involve both the elevation and lateral data (as indeed do many of the preceding parameters). *SΔq* is the root-mean-square slope of the surface, which can be calculated from the same triangular tiling procedure used to generate the visualizations shown earlier. Formally it is defined as the square root of the mean value of the sum of squares of the derivatives of the image in the vertical and horizontal directions, which can be determined simply as the local difference of elevation values between adjacent pixels. The mean summit curvature *Ssc* is similarly related to the second derivatives of elevation, but calculated only for those pixels located at peaks. This depends, of course, on first arriving at a meaningful and accepted definition of which peaks are to be included. The third hybrid property is the ratio of the actual surface area to the projected area *Sdr*. This can be obtained by summing up the areas of the triangles making up the visualization.

None of these hybrid properties is very difficult to compute, but they all depend critically on the sampling interval or spacing of the pixels. Changing lateral resolution will alter the parameter values dramatically so that they are not really functions of the surface but of the measurement technique, and can be used only for comparisons in the most limited way. This is also the case for many of the profile-based measurements, but one of the goals in moving to area-based measurements was to overcome some of the limitations of the older methods. In fact, many engineering surfaces have been shown to have a fractal geometry whose actual surface area is undefined (it increases without limit as the lateral resolution of the measurements improves).

It is a more subtle point, but measurements like these also depend upon whether the elevation data at each pixel are samples of the surface or averages over the pixel area. The mathematics apply for the case of sampling, where the elevation at each pixel is measured at a precise mathematical point and whatever happens between that pixel and the next is not taken into account. In fact, many measurement methods such as conventional stylus instruments and optical interferometers perform some averaging of measurement over the entire area of the pixel, which may either report the maximum value in that area or a weighted average of the elevation values. The mathematics appropriate to these cases has not been worked out and would affect not just the hybrid parameters but all of the parameters described here.

Functional parameters are intended to relate surface geometric data to specific aspects of surface performance, and these are generally related to mechanical engineering applications since the greatest use of surface metrology has thus far been in that field. One typical example is the surface bearing area ratio *S_{tp}*, which is the fraction of the image area that would be in contact with a flat plane parallel to the base if a given height of all peaks was removed by wear (**Figure 45**). This value, of course, can be read directly from the histogram of the elevation data in the image.

Similarly, the amount of volume removed in the process (the material volume ratio *S_{mr}*) can be calculated by integrating the histogram, or using the cumulative histogram. The void volume ratio *S_{vr}* is the volume of empty space within the surface of the specimen, which may be available for retaining or distributing a lubricant. It is measured by integrating the spaces at each elevation level, but this can also be done efficiently using the cumulative histogram. **Figure 46** shows the same data from **Figure 45** but as a surface visualization that reveals the nature of the contact surface and the void volume after some of the peaks have been removed by wear (but assuming that there is no deformation of the remaining surface nor filling in of pits with debris).




Figure 45. Thresholding the range image (in this example, the shotblasted surface from **Figure 1d**) at any particular elevation (in this example 31% below the maximum value) shows the surface area that would be in contact with a plane after a corresponding amount of wear (in the absence of any elastic or plastic deformation).





Figure 46. Visualization of the data from **Figure 45**:

(a) the original shotblasted surface;

(b) the same surface truncated 31% below the top of the highest peak, showing the contact areas and the void volume.

More information is really needed about the contact areas than these parameters provide. The size of the individual contacts is important for heating and deformation, and the void volume may either be completely connected, or consist of isolated pockets, or be a mixture of the two, with very different consequences for lubrication. There are other functional parameters proposed to deal with these and other aspects of surface performance, but these become very specific to each application and will require considerable research to properly define. Many of them are handicapped to a significant degree because the surface elevation data in a range image are single valued. The elevation recorded at each pixel is the maximum height at that point, as detected by a stylus or optical reflection, etc. Undercuts, caves, or pores within the surface that do not show up in the range image may become important if wear removes some of the surface overburden.

Image processing and analysis using the tools already developed in preceding chapters can be used to obtain many of the parameters of interest for surfaces from range images. For example, min and max ranking operators (grey-scale erosion and dilation) can be used to modify the image to form the envelope of the surface which a contact of known form would feel, basically a two-dimensional form of the motif logic mentioned above for profiles. Cross correlation with the image of a defect (crack, dust particle, etc.) can be used to locate such defects. Measurement of features obtained by thresholding can provide data on the contact areas and their distribution after wear has modified a surface. Skeletonization of the pore volume can be used to determine its connectivity

as a pathway to distribute lubricants. Using these tools is straightforward once the significant parameters have been determined so that their relationships to surface behavior and history can be assessed.

New approaches — topographic analysis and fractal dimensions

The limitations and inadequacies of the traditional methods of analysis discussed above prevent them from fully describing real surfaces. They are primarily being used for process control applications in mechanical engineering, where comparison of measurements with prior history provides an indication of change, so long as the measurement technique and instrumentation remains unchanged. Newer methods have become available, but their full meaning and interpretation still remain to be explored. It is hoped that these new approaches that may provide more insight into the description and makeup of surfaces.

Human vision uses global topographic information to organize information on surfaces (Scott, 1995). The arrangement of hills and dales, ridges, courses and saddle points, contains quite a bit of information for describing a surface. A landscape or surface can be divided into regions consisting of hills (points from which all uphill paths lead to one particular peak) and dales (points from which all downhill paths lead to a pit). Boundaries between hills are courses and boundaries between dales are ridge lines (**Figure 47**).

A Pfalz graph (Pfalz, 1976) or change tree (**Figure 48**) connecting the peaks and dales through the respective saddle points where ridge and course lines meet summarizes the topological structure. The change tree can represent directly the height difference and lateral distance between features, which makes decisions straightforward about eliminating features that have either small vertical or lateral extent. This is a direct extension to surfaces of the motif combination used for profiles. Scott (1997) has proposed methods for dealing with the finite extent of real images and the corrections necessary for dealing with the intersection of ridges and courses with the edges of the image area. It is not yet clear just how this information will be used for surface measurement, but parameters such as the volume of connected valleys, the spatial distribution of valleys and peaks across the surface, and orientation of watercourses and ridges seem likely to be important for surface characterization.

At quite a different extreme of local roughness, many surfaces (but emphatically not all) are characterized by a self-similarity (more exactly, a self-affinity) that can be described by a fractal dimension. There are several ways to measure this (which do not exactly agree numerically) plus the need to provide an additional parameter that describes the magnitude of the roughness, and perhaps others to describe the directionality of the surface. The appeal of the fractal dimension is that it is not dependent on the measurement scale, and that it summarizes much of the “roughness” of surfaces in a way that seems to correspond to both the way nature works and the way humans perceive roughness. Given a series of surfaces, the “rougher” the surface as it appears to human

Figure 47. A contour map representing a surface, with the peaks (red), valleys (green) and saddle points (purple) marked.





Figure 48. The Pfalz graph (a) showing which peaks, saddle points, and valleys are connected in **Figure 47**, and the change tree (b) that can be drawn to show the height difference and lateral distance between them.



interpretation (for a variety of basic reasons), the higher the fractal dimension. At the same time it must be noted that the recognition of fractal geometry (even the name) is comparatively new, and there is a “bandwagon” tendency that probably causes it to be applied where it should not be, or with more enthusiasm than critical thinking.

The fractal dimension of a surface is a real number greater than 2 (the topological dimension of the surface) and less than 3 (the topological dimension of the space in which the surface exists). A perfectly smooth surface (dimension 2.0) corresponds to Euclidean geometry, and a plot of the actual area of the measured area as a function of measurement resolution would not change. But for real surfaces an increase in the magnification or resolution with which it is examined will reveal more nooks and crannies and the surface area will increase. For a surprising variety of natural and man-made surfaces, a plot of the area as a function of resolution is linear on a log-log graph, and the slope of this curve gives the dimension D as shown in **Figure 49**.

This description comes directly from the earlier recognition that boundary lines around islands had a length that depended upon the measurement scale. A so-called Richardson plot of the length of the west coast of Britain (Mandelbrot, 1967) as a function of the length of the measurement tool showed this log-log relationship and was one of the triggering ideas that led Mandelbrot to study the mathematics of self-similar structures (ones that appear equally irregular at all scales) and to coin the name fractal for the field. Many other subsequent publications have shown that an extremely broad variety of surfaces also exhibit this kind of geometry, investigated a number of ways to measure the dimension, and begun to study the relationships between the dimension and the history and performance characteristics of surfaces.




Figure 49. Schematic diagram of fractal dimension measurement. As the measurement resolution becomes smaller, the total measured area increases.

Measuring the surface area over a range of resolutions is in fact a rather difficult thing to do (one way is by adsorbing molecules of different sizes), and for basic reasons is not actually appropriate for many surfaces because they are not ideally self-similar. For most surfaces, the lateral directions and the normal direction are distinct in dimension and physical properties, which means that the scaling or self-similarity that exists in one direction may not be the same as the others. At a sufficiently large scale most surfaces approach an ideal Euclidean flat surface. For anisotropic surfaces this situation is more severe and even lateral directions are different. This means that the surfaces are mathematically self-affine rather than self-similar. The fact that elevation measurements are single valued and cannot reveal undercuts means that the measured data would be self-affine even for a truly self-similar surface (for instance one produced by diffusion-limited aggregation of particles on a substrate). For self-affine surfaces and data sets there are still a variety of correct and practical measurement techniques. A few of the more practical ones will be summarized here (and a more complete discussion is available in Russ, 1994).

In most cases, the most robust measure of the fractal dimension is the same procedure that can characterize surfaces that are not ideally fractal, nor perfectly isotropic (Russ, 2001). The Fourier power spectrum can also be used to characterize the instrumental response function, to distinguish it from the surface information. Instead of the usual display mode for the power spectrum, a plot of $\log(\text{magnitude})$ vs. $\log(\text{frequency})$ reveals a fractal surface as a straight line plot whose slope gives the dimension. The principal drawbacks to using the power spectrum plot to measure the dimension are that it tends to overestimate the numerical value of the dimension for relatively smooth surfaces (dimensions between 2.0 and about 2.3), and that the numerical precision of the measured value is lower than some of the other methods can provide for images of a given size. **Figure 50** shows an example of the power spectrum plot for a fractal surface, whose slope gives the fractal dimension.




Figure 50. The Fourier transform power spectrum (upper left) of the shotblasted metal image (**Figure 1d**), with its subsequent analysis. The plot of $\log$ magnitude vs. $\log$ frequency (lower left) averaged over all directions shows a linear relationship which confirms the fractal behavior, and whose slope gives the dimension (2.24). A rose plot of the slope of as a function of orientation (upper left, bold line) shows that the surface is isotropic and has the same dimension in all directions. The thin line on the same plot shows the intercept of the plot as a function of direction, which is a measure of the amplitude of the roughness and also shows isotropy for this surface. Finally, a plot of the distribution of phases of the terms in the Fourier transform shows them to be uniformly random, which is required for a fractal.

Generating the two-dimensional Fourier transform of the surface range image reveals any fractal anisotropy (which can be either weak anisotropy in which the dimension is the same in all directions but the magnitude is not, or strong anisotropy in which the dimension also varies). Plotting the slope and intercept of the plot of $\log(\text{magnitude})$ vs. $\log(\text{frequency})$ as a function of orientation provides a quantitative tool to describe the fractal anisotropy. **Figure 51** shows an example. The separation of the low frequency data that describes the figure and high frequencies that often reveal instrument limitations from the intermediate frequencies can be used to isolate the surface fractal dimension.

By itself, the fractal dimension is only a partial description of surface roughness even for ideally fractal surfaces. Stretching the surface vertically to increase the magnitude of the roughness does not change the slope of the power spectrum or the fractal dimension. An additional measure, which has units of length, is needed to characterize the magnitude. The intercept of the plot of the power spectrum has units of length and can be used for this purpose. So can the topothesy, which is defined as the horizontal distance over which the mean angular change in slope is one radian.

A variety of other measurement approaches can be properly used with self-affine fractal surfaces. Two widely used techniques that deal with the range image of the area directly are the covering blanket and the variogram. Both work correctly for isotropic surfaces but do not reveal anisotropy and can produce nonsense values in those cases rather than an average. The latter is simply a plot of the variance in elevation values as a function of the size of the measured region. Values from small areas placed systematically or randomly over the surface are averaged, and a single mean value obtained. This process is repeated at many different sizes and a plot (**Figure 52**) made that gives the dimension.

The covering blanket or Minkowski method measures the difference (summed over the entire image) between an upper and lower envelope fitted to the surface as a function of the size of the neighborhood used. The minimum and maximum brightness rank operators discussed under image processing are applied with different diameter neighborhoods, and the total difference between them added up. This is analogous to the Minkowski dimension for a profile that was described in Chapter 7, obtained by using the Euclidean distance map to measure the area as a function of distance from the boundary line. The covering blanket method produces a plot as shown in **Figure 53** that gives a dimension. Notice that these three methods give only approximate

Figure 51. The same data presentation as in **Figure 50**, but for the anisotropic ground surface shown in **Figure 1b**. Both the slope and intercept of the power spectrum plot are different in the vertical and horizontal directions. The surface is an anisotropic fractal.




Figure 52. Plot of mean variance vs. neighborhood size to determine the fractal dimension of the shotblasted surface (same image as **Figure 50**).





Figure 53. Plot of the Minkowski cover volume as a function of neighborhood size to determine the fractal dimension (same image as **Figures 50** and **52**).

agreement as to the numerical value of the dimension. Part of this is just finite measurement precision, but part of the difference arises from the fact that all of these techniques measure something that is slightly different. These are limits to the actual dimension and in general will not agree, so when comparisons are being made between surfaces it is important to always use one technique for all of the measurements.

It is often attractive to perform measurements in a lower dimension, since a smaller number of data points are involved. Historically, much of the work with fractal measurement has been done with boundary lines, whose dimension lies between 1.0 (the Euclidean or topological dimension of a line) and 1.999 (a line whose irregularity is so great that it wanders across an entire plane). There is a way to do this with fractal surfaces, by intersecting the surface with a plane and then measuring the dimension of the line that is the intersection. It is vitally important however that this plane be parallel to the nominal surface orientation rather than a vertical cut. The vertical cut would produce the same profile as that obtained with a profilometer, but because the surface is self-affine and not self-similar the proper measurement of this profile is complicated and the common techniques such as the Richardson plot mentioned above do not apply. Also, of course, the profile will be oriented in a particular direction and cannot be used with anisotropic surfaces.

The horizontal cut is called a slit-island method, and it corresponds exactly to the case with which Richardson was dealing. The horizontal plane corresponds to sea level and the outlines are the coastlines of the islands produced by the hills that rise above the sea. A plot of the length of these coastlines as a function of measurement scale produces a dimension that is exactly one less than the surface dimension (the difference between the topological dimensions of a surface and a line). It is not usually convenient to measure the length of a coastline using a computer in the way Richardson did, by setting a pair of dividers to a particular scale length and "striding" around the





Figure 54. Contour lines produced by thresholding the image from **Figure 1d** at 50% of its elevation range and the plot from a box-counting measurement.

coastline so that the boundary length was the divider setting times the number of strides; but a variety of other methods are readily implemented in a computer.

The analog to the Minkowski blanket for the surface is a Minkowski sausage produced by thickening the boundary line by various amounts and plotting the area covered vs. the width of the stripe or sausage. The Euclidean distance map discussed in Chapter 7 on binary image processing accomplishes this procedure very efficiently and without the directional variation that results from conventional pixel-based dilation. Another method is box-counting (**Figure 54**) in which a grid is placed on the image and the number of squares through which the boundary line passes is counted as a function of the size of the grids. While these methods are quite fast, they are only applicable to isotropic surfaces.

It is beyond the scope of this text to discuss the relationships that have been found between fractal dimensions and various aspects of the history and properties of surfaces. In brief, most processes of surface formation that involve brittle fracture or deposit large amounts of energy in small regions tend to produce fractal surfaces, and the numerical value of the dimension is often a signature of the process that was involved (Russ, 1997). Likewise, many surface contact applications (electrical, thermal, etc.) depend upon the relationship between contact area and pressure and fractal geometry is pertinent to this case. Under some circumstances friction and wear may also be related to the surface dimension. Fractal description of surfaces is a new, and somewhat "trendy" approach to surface characterization. Although it clearly offers powerful methods that apply to some surfaces, for others such as ductile deformation it is inappropriate.

References

- M. Abbasi-Dezfouli, T. G. Freeman (1994). Patch matching in stereo-images based on shape, *ISPRS Int. Arch. of Photogramm. and Remote Sensing* 30(3/1):1-8.
- M. R. Anderberg (1973). *Cluster Analysis for Applications*. Academic Press, New York.
- Agfa Compugraphic Division (1992). *Digital Color Prepress* (vols. 1 & 2). Agfa Corp., Wilmington, MA.
- E. H. Aigeltinger, K. R. Craig, R. T. DeHoff (1972). Experimental determination of the topological properties of three-dimensional microstructures. *J. Microsc.* 95:69-81.
- G. R. Arce, J. L. Paredes, J. Mullan (2000). Nonlinear filtering for image analysis and enhancement, in *Handbook of Image and Video Processing* (A. Bovik, Ed.). Academic Press, San Diego.
- J. Astola, P. Haavisto, Y. Neuvo (1990). Vector median filters. *Proc. IEEE* 78:678-689.
- N. Baba, M. Naka, Y. Muranaka, et al. (1984). Computer-aided stereographic representation of an object reconstructed from micrographs of serial thin sections. *Micron and Microsc. Acta* 15:221-226.
- N. Baba, M. Baba, M. Imamura, et al. (1989). Serial section reconstruction using a computer graphics system: application to intracellular structures in yeast cells and to the periodontal structure of dogs' teeth. *J. Electron Microsc. Tech.* 11:16-26.
- A. J. Baddeley, H. J. G. Gundersen, L. M. Cruz-Orive (1986). Estimation of surface area from vertical sections. *J. Microsc.* 142:259-276.
- A. J. Baddeley, C. V. Howard, A. Boyde, et al. (1987). Three-dimensional analysis of the spatial distribution of particles using the tandem-scanning reflected light microscope. *Acta Stereol.* 6 (Suppl. II) 87-100.
- R. Balasubramanian, C. A. Bouman, J. P. Allebach (1994). Sequential scalar quantization of color images. *J. Electron. Imaging* 3(1):45-59.
- D. Ballard (1981). Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition* 13(2):111-122.
- D. H. Ballard, C. M. Brown (1982). *Computer Vision*. Prentice Hall, Englewood Cliffs, NJ.
- D. H. Ballard, C. M. Brown, J. A. Feldman (1984). An approach to knowledge-directed image analysis, in *Computer Vision Systems* (A. R. Hanson, E. M. Riseman, Eds.) Academic Press, New York, NY, 271-282.
- S. T. Barnard, W. B. Thompson (1980). Disparity Analysis of Images. *IEEE Trans. on Patt. Anal. Mach. Intell.* (PAMI) 2(4):333-340.

- F. L. Barnes, S. G. Azavedo, H. E. Martz, G. P. Roberson, D. J. Schneberk, M. F. Skeate (1990). *Geometric Effects in Tomographic Reconstruction*. Lawrence Livermore National Laboratory Rept UCRL-ID-105130.
- H. G. Barth, S-T. Sun (1985) Particle Size Analysis. *Anal. Chem.* 57:151.
- A. Barty, K. A. Nugent, A. Roberts, et al. (2000). Quantitative Phase Tomography. *Optics Communications* 175:329-336.
- J. K. Beddow, G. C. Philip, A. F. Vetter (1977). On relating some particle profiles characteristics to the profile Fourier coefficients. *Powder Technol.* 18:15-19.
- G. Bertrand, J-C. Everat, M. Couprie (1997). Image segmentation through operators based on topology. *J. Electron. Imaging* 6(4):395-405.
- V. Berzins (1984) Accuracy of Laplacian edge detectors. *Comput. Vis. Graph. Image Proc.* 27:1955-2010.
- P. J. Besl, N. D. McKay (1992). A method for registration of 3-D shapes. *IEEE Trans. Patt. Anal. Mach. Intell.* 14(2): 239-256.
- S. Beucher, C. Lantjeoul (1979). Use of Watersheds in Contour Detection. *Proc. Intl. Workshop Image Process*, CCETT, Rennes, France.
- G. B. Binnig, C. F. Quate, C. Gerber (1986). Atomic Force Microscope. *Phys. Rev. Lettr.* 52:930.
- J. L. Bird, D. T. Eppler, D. M. Checkley, Jr. (1986). Comparisons of herring otoliths using Fourier series shape analysis. *Can. J. Fish. Aquat. Sci.* 43:1228-1234.
- W. J. Black (1986). *Intelligent Knowledge-Based Systems: An Introduction*. Van Nostrand Reinhold, London.
- R. A. S. Blackie, R. Bagby, L. Wright, J. Drinkwater, S. Hart (1987). Reconstruction of three-dimensional images of microscopic objects using holography. *Proc. R. Microsc. Soc.* 22:98.
- F. Boddeke (1998). *Quantitative Fluorescence Microscopy. ASCI Dissertation Series*. Delft Univ. Press, Delft, Netherlands.
- F. R. Boddeke, L. J. van Vliet, H. Netten, I. T. Young (1994). Autofocusing in microscopy based on the OTF and sampling. *Bioimaging* 2:193-203.
- G. Borgefors (1996). On digital distance transforms in three dimensions. *Comp. Vision Image Understanding* 64(3):368-376.
- G. Borgefors, I. Nyström, G. Sanniti di Baja (1999). Computing skeletons in three dimensions. *Patt. Recognition* 32(7):1225-1236.
- S-T. Bow (1992). *Pattern Recognition and Image Preprocessing*. Marcel Dekker, New York.
- A. Boyde (1973). Quantitative photogrammetric analysis and quantitative stereoscopic analysis of SEM images. *J. Microsc.* 98:452.
- S. Bradbury, B. Bracegirdle (1998). *Introduction to Light Microscopy*. Bios Scientific Publishers, Oxford, U.K.
- B. Bracegirdle, S. Bradbury (1995). *Modern Photomicrography*. Bios Scientific Publishers, Oxford UK.
- R. N. Bracewell (1984). The fast Hartley transform. *Proc. IEEE* 72:8.
- R. N. Bracewell (1986). *The Hartley Transform*. Oxford Univ. Press, Oxford, U.K.
- R. N. Bracewell (June 1989). The Fourier Transform. *Scientific American*,
- G. Braudaway (1987). A procedure for optimum choice of a small number of colors from a large color palette for color imaging. *Proc. Electronic Imaging 1987*, San Francisco, CA.
- L. G. Briarty, P. H. Jenkins (1984). GRIDSS: an integrated suite of microcomputer programs for three-dimensional graphical reconstruction from serial sections. *J. Microsc.* 134:121-124.

- D. S. Bright, D. E. Newbury, R. B. Marinenko (1988). Concentration-concentration histograms: scatter diagrams applied to quantitative compositional maps, in *Microbeam Analysis 1988* (D. E. Newbury Ed.). San Francisco Press, 18–24.
- P. Brodatz (1966). *Textures: A Photographic Album for Artists and Designers*, Dover, New York.
- L. G. Brown (1992). A survey of image registration techniques. *ACM Computing Surveys* 24(4):325–376.
- R. K. Bryan, J. Skilling (1980). Deconvolution by maximum entropy, as illustrated by application to the jet of M87. *Mon. Not. R. Ast. Soc.* 191: 69–79.
- V. Buzuloiu, M. Ciuc, R. M. Rangayyan, et al. (2001). Adaptive-neighborhood histogram equalization of color images. *J. Electron. Imaging* 10(2):445–459.
- I. J. Canny (1986). A computational approach to edge detection. *IEEE Trans. Patt. Anal. Mach. Intell.* (PAMI) 8(6):679–698.
- C. Carlsen (1985). Reconstruction of true surface topographies in scanning electron microscopes using backscattered electrons. *Scanning* 7:169–177.
- W. A. Carrington (1990). Image restoration in 3D microscopy with limited data, in *Bioimaging and Two Dimensional Spectroscopy*, *Proc. SPIE*, Vol. 1205 (L. C. Smith, Ed.), 72–83.
- E. Castle, P. A. Zhdan (1997). Characterization of surface topography by SEM and SFM: Problems and solutions. *J. Phys. D. App. Phys.* 30:722.
- K. R. Castleman (1979). *Digital Image Processing* Prentice Hall, Englewood Cliffs, NJ.
- R. L. T. Cederberg (1979). Chain-link coding and segmentation for raster scan devices. *Comput. Graph. Image Process* 10:224–234.
- P. Chieco, A. Jonker, C. Melchiorri, et al. (1994). A user's guide for avoiding errors in absorbance image cytometry. *Histochemical J.* 26:1–19.
- C. K. Chui (1992). *An Introduction to Wavelets*. Academic Press, London.
- J. R. B. Cockett (1987). Decision expression optimization. *Fundamenta Informaticae X* 93–114.
- L. Cohen (1991). On active contour models and balloons. *Comp. Vision, Graph. Image Process.* 53(2):211–218.
- M. L. Comer, E. J. Delp (1999). Morphological operations for color image processing, *J. Electron. Imaging* 8:279–289.
- J. Condeco, L. H. Christensen, S. F. Jorgensen, J. C. Russ, B-G. Rosen (2000). A comparative study of image stitching algorithms for surface topography measurements. *X. Intl Colloq. on Surfaces*, Germany.
- J. Cookson (1994). Three-dimensional reconstruction in microscopy. *Proc. R. Microscopical Soc.* 29(1):3–10.
- J. W. Cooley, J. W. Tukey (1965). An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*.
- A. M. Cormack (1963). Representation of a function by its line integrals with some radiological applications. *J. Appl. Phys.* 34:2722–2727.
- A. M. Cormack (1964). Representation of a function by its line integrals with some radiological applications II. *J. Appl. Phys.* 35:2908–2913.
- T. N. Cornsweet (1970). *Visual Perception*, Academic Press, New York.
- L. F. Costa, R. M. Cesar (2001). *Shape Analysis and Classification*. CRC Press, Boca Raton, FL.
- M. Coster, J.-L. Chermant (1985). *Précis D'Analse D'Images Éditions du Centre National de la Recherche Scientifique*, Paris.
- L-M. Cruz-Orive (1976). Particle size-shape distributions: the general spheroid problem. *J. Microsc.* 107:235; 112:153.

- L-M. Cruz-Orive (1983). Distribution-free estimation of sphere size distributions from slabs showing overprojections and truncation, with a review of previous methods. *J. Microsc.* 131:265.
- P. E. Danielsson (1980). Euclidean distance mapping. *Comput. Graph. Image Process* 14:227–248.
- I. Daubechies (1992). Ten Lectures on Wavelets. *CBMS-NSF Regional Conf. Series Applied Math*, Philadelphia, PA.
- I. Daubechies (1996). Where do wavelets come from? — A personal point of view. *Proc. IEEE* 84(4):510–513.
- D. G. Daut, D. Zhao, J. Wu (1993). Double predictor differential pulse coded modulation algorithm for image data compression. *Optical Engineering* 32(7):1514–1523.
- J. Davidson (1991). Thinning and skeletonization: a tutorial and overview, in *Digital Image Processing: Fundamentals and Applications* (E. Dougherty, Ed.). Marcel Dekker, New York.
- E. R. Davies (1988). On the noise suppression and image enhancement characteristics of the median, truncated median and mode filters, *Patt. Recog. Lett.* 7:87–97.
- H. W. Deckman, K. L. D'Amico, J. H. Dunsmuir, B. P. Flannery, S. M. Gruner (1989). *Microtomography Detector Design Advances in X-ray Anal.* 32:641.
- R. T. Dehoff, F. N. Rhines (1968). *Quantitative Microscopy*, McGraw-Hill, New York.
- D. DeMandolx, J. Davoust (1997). Multicolor analysis and local image correlation in confocal microscopy. *J. Microsc.* 185:21–36.
- M. De Marsicoi, L. Cinque, S. Levialdi (1997). Indexing pictorial documents by their content: a survey of current techniques. *Image Vision Computing* 15:119–141.
- P. A. Devijver, J. Kittler (1980). *Pattern Recognition: A Statistical Approach*. Prentice Hall, Englewood Cliffs, NJ.
- G. Diaz, A. Zuccarelli, I. Pelligra, et al. (1989). Elliptic Fourier analysis of cell and nuclear shapes. *Comput. Biomed. Res.* 22: 405–414. .
- G. Diaz, D. Quacci, C. Dell'Orbo (1990). Recognition of cell surface modulation by elliptic Fourier analysis. *Comput. Meth. Prog. Biomedicine* 31:57–62.
- M. Dietzsch, K. Papenfuss, T. Hartmann (1997). The MOTIF method (ISO 12085) — a suitable description for functional manufactural and metrological requirements, in *7th Intl Conf. on Metrology and Properties of Engineering Surfaces* (B. G. Rosen, R. J. Crafoord, Eds.). Chalmers Univ., Göteborg, Sweden, 231–238.
- E. R. Dougherty, J. Astola (1994). *An Introduction to Nonlinear Image Processing*. SPIE, Bellingham, WA.
- E. R. Dougherty, J. Astola (1999). *Nonlinear Filters for Image Processing*. SPIE, Bellingham, WA.
- N. Draper, H. Smith (1981). *Applied Regression Analysis* (2nd edition). John Wiley & Sons, New York.
- R. O. Duda, P. E. Hart (1972). Use of the Hough transform to detect lines and curves in pictures. *Commun. ACM* 15(1):11–15.
- R. O. Duda, P. E. Hart (1973). *Pattern Classification and Scene Analysis*. John Wiley & Sons, New York, NY.
- T. R. Edwards (1982). Two-dimensional convolute integers for analytical instrumentation. *Anal. Chem.* 54:1519–1524.
- R. Ehrlich, B. Weinberg (1970). An exact method for characterization of grain shape. *J. Sediment. Petrol.* 40:205–212.
- R. Ehrlich, S. K. Kennedy, S. J. Crabtree, R. L. Cannon (1984). Petrographic image analysis: 1. analysis of reservoir pore complexes. *J. Sediment Petrol.* 54:1365–1378.

- A. Engel, A. Massalski (1984). 3D reconstruction from electron micrographs: its potential and practical limitations. *Ultramicroscopy* 13:71-84.
- Y. Fahmy, J. C. Russ, C. Koch (1991). Application of fractal geometry measurements to the evaluation of fracture toughness of brittle intermetallics. *J. Mater. Res.* 6(9):1856-1861.
- C. Faloutsos, W. Equitz, M. Flickner, et al. (1994). Efficient and effective querying by image content. *J. Intelligent Inf. Sys.* 3:231-262.
- J. Feder (1988). *Fractals*. Plenum Press, New York NY.
- L. A. Feldkamp, L. C. Davis, J. W. Kress (1984). Practical cone beam algorithm. *J. Opt. Soc. Am.* 1(6):612.
- S. F. Ferson, F. J. Rohlf, R.K. Koehn (1985). Measuring shape variation of two-dimensional outlines. *Systematic Zoology* 34:59-68.
- L. F. Firestone, K. Cook, et al. (1991). Comparison of autofocus methods for automated microscopy. *Cytometry* 12:195-206.
- Y. Fisher, E.W. Jacobs, R. D. Boss (1992). Fractal image compression using iterated transforms, in *Image and Text Compression*. J. A. Storer, Ed. Kluwer Academic Publishers, Boston, MA, 35-61.
- M. Flickner, H. Sawhney, W. Niblack, et al. (1995). Query by image and video content: the QBIC system, *IEEE Comput.* 28(9):23-32.
- A. G. Flook (1978). Use of dilation logic on the Quantimet to achieve fractal dimension characterization of texture and structured profiles. *Powder Techn.* 21:295-298.
- A. G. Flook (1982). Fourier analysis of particle shape, in *Particle Size Anal.* 1981-2 (N. G. Stanley-Wood, T. Allen, Eds.). Wiley Heyden, London.
- J. D. Foley, A. Van Dam (1984). *Fundamentals of Interactive Computer Graphics* Addison Wesley, Reading, MA.
- J. Frank, Ed. (1992). *Electron Tomography*. Plenum Press, NY.
- R. J. Frank, T. J. Grabowski, H. Damasio (1995). Voxelwise percentage tissue segmentation of human brain magnetic resonance images, in *Abstracts, 25th Annual Meeting, Soc. for Neuroscience*, Washington, 694. Cited in M. Sonka, V. Hlavac, R. Boyle (1999). *Image Processing, Analysis and Machine Vision*. Brooks Cole, Pacific Grove, CA, 132.
- M. Frederik, A. Collignon, D. Vandermeulen, et al. (1997). Multi-modality image registration by maximization of mutual Information, *IEEE Trans. Med. Imaging*, 16(2):187-198.
- H. Freeman (1961). On the encoding of arbitrary geometric configurations. *IEEE Trans.* EC-10:260-268.
- H. Freeman (1974). Computer processing of line-drawing images. *Comput. Surveys* 6:57-97.
- H. Freeman, L. S. Davis (1977). A corner finding algorithm for chain-code curves. *IEEE Trans Comput.* 26:297-303.
- W. Frei, C. C. Chen (1977). Fast boundary detection: a generalization and a new algorithm. *IEEE Trans. Comput.* C-26:988-998.
- B. R. Frieden (1988). A comparison of maximum entropy, maximum a posteriori and median window restoration algorithms, in *Scanning Microscopy* (Suppl. 2) (P. Hawkes, et al., Eds.), Scanning Microscopy International, Chicago, 107-111.
- J. P. Frisby (1980). *Vision: Illusion, Brain and Mind*. Oxford Univ. Press, Oxford, U.K.
- K. S. Fu (1974). *Syntactic Methods in Pattern Recognition*, Academic Press, Boston.
- K. S. Fu, J. K. Mui (1981). A survey on image segmentation. *Pattern Recognition* 13:3-16.
- K. S. Fu (1982). *Syntactic Pattern Recognition and Applications*, Prentice Hall, Englewood Cliffs, NJ.
- H. Fuchs, S. M. Pizer, L. C. Tsai, et al. (1982). Adding a true 3-D display to a raster graphics system. *IEEE Comput. Graphics Applic.* 2:73-78.
- K. Fukunaga (1990). *Statistical Pattern Recognition*, 2nd edition. Academic Press, Boston.

- R. S. Gentile, J. P. Allebach, E. Walowit (1990). Quantization of color images based on uniform color spaces. *J. Imaging Technol.* 16(1):12-21.
- R. C. Gonzalez, P. Wintz (1987). *Digital Image Processing*, 2nd edition. Addison-Wesley, Reading, MA.
- R. C. Gonzalez, R. E. Woods (1993). *Digital Image Processing*, Addison-Wesley, Reading, MA.
- R. Gordon (1974). A tutorial on ART (algebraic reconstruction techniques). *IEEE Trans.* NS-21:78-93.
- Y. A. Gowayed, J. C. Russ (1991). Geometric characterization of textile composite preforms using image analysis techniques, *J. Comput. Assisted Microsc.* 3(#4):189-200.
- G. H. Granlund (1972). Fourier Preprocessing for Hand Print Character Recognition, *IEEE Trans. Comput.* C21(2):195-201.
- A. Grasselli (Ed.). (1969). *Automatic Interpretation and Classification of Images*. Academic Press, Boston.
- C. Gratin, F. Meyer (1992). Morphological three-dimensional analysis. *Scanning Microsc. Suppl.* 6:129-135.
- F. C. A. Green, I. T. Young, G. A. Lighthart (1985). A comparison of different focus functions for use in autofocus algorithms. *Cytometry* 6:81091.
- S. Grossberg (Ed.). (1988). *Neural Computers and Natural Intelligence*. MIT Press, Cambridge, MA.
- P. Gualtieri, P. Coltell (1992). An automated system for the analysis of moving images. *J. Comput. Assisted Microsc.* 3(#1):15-22.
- H. J. G. Gundersen (1986). Stereology of arbitrary particles *J. Microsc.* 143:3-45.
- H. J. G. Gundersen et al. (1988). Some new, simple and efficient stereological methods and their use in pathological research and diagnosis. *Acta Pathologica, Microbiologica et Immunologica Scandinavica* 96:857.
- H. J. G. Gundersen, E. B. Jensen (1987). The efficiency of systematic sampling in stereology and its prediction. *J. Microsc.* 147:229-263.
- D-P. Hader, Ed. (1992). *Image Analysis in Biology*, CRC Press, Boca Raton, FL.
- K. J. Halford, K. Preston (1984). 3-D Skeletonization of Elongated Solids. *Computer Vision Graphics Image Process.* 27:78-91.
- Y. S. Ham (1993). Differential absorption cone-beam microtomography, Ph.D. Thesis, North Carolina State University,
- D. J. Hand (1981). *Discrimination and Classification*, John Wiley & Sons, New York, NY.
- R. M. Haralick (1978). Statistical and structural approaches to texture, *Proc. 4th Intl. Joint Conf. Patt Recog.*, Kyoto, Japan, 45.
- R. Haralick (1979). Statistical and textural approaches to textures. *Proc. IEEE* 67:786-804.
- R. M. Haralick, I. Dinstein (1975). A spatial clustering procedure for multi-image data. *Comput. Graph. Image. Proc.* 12:60-73.
- R. M. Haralick, L. G. Shapiro (1988). Segmentation and its place in machine vision. *Scanning Microscopy Supplement* 2:39-54.
- R. M. Haralick, L. G. Shapiro (1992). *Computer and Robot Vision I*. Addison-Wesley, Reading, MA.
- R. M. Haralick, K. Shanmugam, I. Dinstein (1973). Textural features for image classification. *IEEE Trans. SMC-3*:610-621.
- J. A. Hartigan (1975). *Clustering Algorithms*. John Wiley & Sons, New York.
- R. V. L. Hartley (1942). A more symmetrical Fourier analysis applied to transmission problems. *Proc. IRE*.
- S. Haykin (1993). *Neural Networks*. MacMillan, New York.

- D. Hearn, M. P. Baker (1986). *Computer Graphics*. Prentice Hall, Englewood Cliffs, NJ.
- M. Heath, S. Sarkar, T. Sanocki, et al. (1997). A robust visual method for assessing the relative performance of edge-detection algorithms. *IEEE Trans. Patt. Anal. Mach. Intell.* (PAMI)19(12):1338–1359.
- D. O. Hebb (1949). *The organization of behaviour*. John Wiley & Sons, New York.
- P. Heckbert (1982). Color image quantization for frame buffer display. *Computer Graphics* 16(3):297–307.
- R. Hegerl (1989). Three-dimensional reconstruction from projections in electron microscopy. *European J. Cell Biol.* 48 (Suppl. 25):135–138.
- H. Heijmans (1991). Theoretical aspects of grey-level morphology. *IEEE Trans. Patt. Anal. Mach. Intell.* (PAMI) 13(6):568–582.
- H. Heijmans (1994). *Morphological Image Operators*. Academic Press, New York.
- C. Heipke (1992). A Global Approach for Least-Squares Image Matching and Surface Reconstruction in Object Space. *Photogrammetric Eng. Remote Sensing* 58(3):317–323.
- G. T. Herman (1980). *Image Reconstruction from Projections — The Fundamentals of Computerized Tomography*. Academic Press, New York.
- E. C. Hildreth (1983). The detection of intensity changes by computer and biological vision systems. *Comput. Vis. Graph. Image Proc.* 22:1–27.
- T. J. Holmes, S. Bhattacharyya, J. A. Cooper, et al. (1995). Light microscopic images reconstructed by maximum likelihood deconvolution, in *Handbook of Biological Confocal Microscopy* (J. Pawley, Ed.). Plenum Press, New York, 389–402.
- T. J. Holmes, N. J. O'Connor (2000). Blind deconvolution of 3D transmitted light brightfield micrographs. *J. Microsc.* 200(2):114–127.
- B. J. Holt, L. Hartwick (1994). Visual image retrieval for applications in art and art history. *Proc. Storage and Retrieval for Image and Video Databases II*, SPIE, San Jose.
- B. K. P. Horn (1970). *Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View* (AI Tech Report 79, Project MAC). Mass. Inst. Tech., Cambridge, MA.
- B. K. P. Horn (1975). Obtaining shape from shading information, in *Psychology of Computer Vision* (P. H. Winston, Ed.). McGraw Hill New York, 115–155.
- P. Hough (1962). Method and means for recognizing complex patterns, U.S. Patent 3,069,654.
- C. V. Howard, M. G. Reed (1998). *Unbiased Stereology, Three Dimensional Measurements in Stereology*. BIOS Scientific Publishers, Oxford, U.K.
- J. Huang, S. M. Dunn, S. M. Wiener, et al. (1994). A Method for Detecting Correspondences in Stereo Pairs of Electron Micrographs of Networks. *J. Comp. Assisted Microsc.* 6(2):85–102.
- T. S. Huang (1979). A fast two-dimensional median filtering algorithm. *IEEE Trans. ASSP* 27:13–18.
- D. H. Hubel (1988). *Eye, Brain, and Vision*, Scientific American Library. W. H. Freeman, New York.
- H. E. Hurst, R. P. Black, Y. M. Simaika (1965). *Long-Term Storage: An Experimental Study*. Constable, London.
- C. Hwang, S. Venkatraman, K. R. Rao (1993). Human visual system weighted progressive image transmission using lapped orthogonal transform classified vector quantization. *Optical Eng.* 32(7):1524–1530.
- S. Inoue (1986). *Video Microscopy*, Plenum Press, New York.
- B. Jahne (1997). *Practical Handbook on Image Processing for Scientific Applications*. CRC Press, Boca Raton, FL.
- A. K. Jain (1989). *Fundamentals of Digital Image Processing*. Prentice Hall, London.

- M. James (1988). *Pattern Recognition*. Blackwell Scientific, London.
- J. R. Janesick (2001). *Scientific Charge-Coupled Devices*. SPIE Press, Bellingham, WA.
- E. T. Jaynes (1985). Where do we go from here? in *Maximum Entropy and Bayesian Methods in Inverse Problems* (C. R. Smith, W. T. Grandy, Eds.). D. Reidel Publishing Co., Dordrecht, Holland, 21–58.
- K. F. Jarausch, T. J. Stark, P. E. Russell (1996). Silicon structures for in-situ characterization of atomic force microscopy probe geometry. *J. Vac. Sci. Technol. B*. 14(6):3425.
- J. P. Jernot (1982). Thèse de Doctorat ès Science, Université de Caen, France.
- E. M. Johnson, J. J. Capowski (1985). Principles of reconstruction and three-dimensional display of serial sections using a computer, in *The Microcomputer in Cell and Neurobiology Research* (R. R. Mize, Ed.). Elsevier, New York, 249–263.
- L. R. Johnson, A. K. Jain (1981). An Efficient Two-Dimensional FFT Algorithm. *IEEE Trans. PAMI* 3:698–701.
- Q. C. Johnson, J. H. Kinney, U. Bonse, et al. (1986). Micro-Tomography using Synchrotron Radiation Lawrence Livermore National Laboratory Preprint UCRL-93538 A. C. Kak, M. Slaney (1988). *Principles of Computerized Tomographic Imaging*. IEEE Pub. PC-02071.
- S. Joshi, M. I. Miller (1993). Maximum a posteriori estimation with good roughness for 3D optical sectioning microscopy. *Opt. Soc. Am. A*. 10:1078–1985.
- M. Kaes., A. Witkin, D. Terzopoulos (1987). Snakes: active contour models, *J. Comput. Vision* 1:321–331.
- A. C. Kak, M. Slaney (1987, 2001). *Principles of Computerized Tomographic Imaging*. SIAM, Philadelphia PA.
- H. R. Kang (1997). *Color Technology for Electronic Imaging Devices*. SPIE Optical Engineering Press, Bellingham, WA.
- H. R. Kang (1999). *Digital Color Halftoning*. SPIE Optical Engineering Press, Bellingham, WA.
- J. N. Kanpur, P. K. Sahoo, A. K. C. Wong (1985). A new method for grey-level picture thresholding using the entropy of the histogram. *Comput. Vision, Graph. Image Process.* 29:273–285.
- N. Karssemeijer et al. (Eds.). *Digital Mammography*. Kluwer Academic Publishers, New York.
- M. Kass, A. Witkin, D. Terzopoulos (1987). Snakes: active contour models. *Int. J. Comput. Vision* 1(4):321–331.
- A. E. Kayaalp, R. C. Jain (1987). Using SEM stereo to extract semiconductor wafer pattern topography. *Proc. SPIE* 775:18–26.
- B. H. Kaye, J. E. LeBlanc, G. Clark (1983). A study of physical significance of three-dimensional signature waveforms. *Proc. Fineparticle Characterization Conf.*
- B. H. Kaye (1986). Image analysis procedures for characterizing the fractal dimension of fineparticles. *Proc. Particle Technol. Conf.*, Nürnberg, Germany.
- A. Keating (1993). Duke University, Durham, NC, private communication.
- D. J. Keller (1991). Reconstruction of STM and AFM images distorted by finite-size tips. *Surf. Sci.* 253:353–364.
- D. J. Keller, F. S. Franke (1993). Envelope reconstruction of probe microscope images. *Surf. Sci.* 294:409–419.
- S. Kim, J. Lee, J. Kim (1988). A New Chain-Coding Algorithm for Binary Images Using Run-Length Codes. *Comput. Vision, Graph. Image Process.* 41:114–128.
- J. H. Kinney, Q. C. Johnson, M. C. Nichols, et al. (1989). X-ray microtomography on beamline X at SSRL. *Rev. Sci. Instrum.* 60(7):2471–4.

- J. H. Kinney, M. C. Nichols, U. Bonse, et al. (1990). Nondestructive imaging of materials microstructures using X-ray tomographic microscopy. *Proc. MRS Symp. Tomographic Imaging*. Boston, MA.
- R. Kirsch (1971). Computer Determination of the Constituent Structure of Biological Images. *Comput. Biomed. Res.* 4:315–328.
- J. Kittler, J. Illingworth, J. Foglein (1985). Threshold selection based on a simple image statistic. *Comput. Vision, Graph. Image Process.* 30:125–147.
- V. Kober, M. Mozerov, J. Alvarez-Borrego (2001). Nonlinear filters with spatially connected neighborhood. *Opt. Eng.* 40(6):971–983.
- A. Kriete, Ed. (1992). *Visualization in Biomedical Microscopies: 3D Imaging and Computer Applications*, VCH Publishers, Weinheim.
- F. P. Kuhl, C. R. Giardina (1982). Elliptic Fourier features of a closed contour. *Comp. Graph. Image Process.* 18:236–258.
- K. J. Kurzydowski, B. Ralph (1995). *The Quantitative Description of the Microstructure of Materials*, CRC Press, Boca Raton, FL.
- M. Kuwahara, K. Hachimura, S. Eiho, M. Kinoshita (1976). Processing of RI-angiocardigraphic images, in *Digital Processing of Biomedical Images* (K. Preston, M. Onoe, Eds.), Plenum Press, New York, 187–202.
- R. L. Lagendijk, J. Biemond (1991). *Iterative identification and restoration of images*. Kluwer Academic Publishers, Boston, MA.
- L. Lam, S. Lee, C. Suen (1992). Thinning methodologies — a comprehensive survey. *IEEE Trans. PAMI* 14:868–885.
- C. Lantéjoul, S. Beucher (1981). On the use of the geodesic metric in image analysis. *J. Microsc.* 121:39.
- R. S. Ledley, M. Buas, T. J. Golab (1990). Fundamentals of true-color image processing. *Proc. Intl. Conf. Patt. Recog.* 1:791–795.
- D. L. Lee, A. T. Winslow (1993). Performance of three image-quality metrics in ink-jet printing of plain papers. *J. Electron. Imaging* 2(3):174–184.
- S. U. Lee, S. Y. Chung, R. H. Park (1990). A comparative performance study of several global thresholding techniques for segmentation. *Comput. Vision. Graph. Image Process.* 52:171–190.
- P. E. Lestrel, Ed. (1997). *Fourier Descriptors and Their Applications in Biology*, Cambridge Univ. Press, Cambridge, MA.
- J. Y. Lettvin, R. R. Maturana, W. S. McCulloch, W. H. Pitts (1959). What the frog's eye tells the frog's brain. *Proc. Inst. Rad. Eng.* 47(#11), 1940–1951.
- S. Levialdi (1972). On shrinking binary picture patterns. *Commun. ACM* 15(1):7–10.
- H. Li, M. Novak, R. Forchheimer (1993). Fractal-based image sequence compression scheme. *Optical Eng.* 32(7):1588–1595.
- B. Lichtenbelt, R. Crane, S. Naqvi (1998). *Introduction to Volume Rendering*. Prentice Hall, Saddle River, NJ.
- W. Lin, T. J. Holmes, D. H. Szarowski, J. N. Turner (1994). Data corrections for three-dimensional light microscopy stereo pair reconstruction. *J. Comput. Assisted Microsc.* 6(3):113–128.
- M. Lineberry (1982). Image segmentation by edge tracing. *Appl. Digital Image Process.* IV, Vol. 359.
- S. Lobregt, P. W. Verbeek, F. C. A. Groen (1980). Three-dimensional skeletonization: principle and algorithm. *IEEE Trans. PAMI* 2:75–77.
- E. Mach (1906). Über den Einfluss räumlich und zeitlich variierender Lichtreize auf die Gesichtswahrnehmung. *S.-B. Akad. Wiss. Wien, Math.-Nat. Kl.* 115:633–648.

- J. B. MacQueen (1967). Some Methods for the Classification and Analysis of Multivariate Observations, *Proc. 5th Berkeley Symp. Mathematical Statistics Probab.*, 1, 281–297.
- E. Mainsah, K. J. Stout, T. R. Thomas (2001). Surface measurement and characterization in *Metrology and Properties of Engineering Surfaces* (E. Mainsah, Ed.). Kluwer Academic Publishers, London, 1–42.
- S. G. Mallat (1989). A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. PAMI* 11(7):674–693.
- T. Malzbender, D. Gelb, D. Wolters (2001). Polynomial texture maps. *Comput. Graph., Proc. Siggraph 2001*, pp.519–528.
- B. B. Mandelbrot (1967). How long is the coast of Britain? Statistical self-similarity and fractional dimension. *Science* 155:636–638.
- B. B. Mandelbrot (1982). *The Fractal Geometry of Nature*. W. H. Freeman, San Francisco, CA.
- B. B. Mandelbrot, D. E. Passoja, A. J. Paullay (1984). Fractal character of fracture surfaces of metals. *Nature* 308:721.
- R. Mann, S. Stanley, D. Vlaev, et al. (2001). Augmented reality visualization of fluid mixing in stirred chemical reactors using electrical resistance tomography *J. Electron. Imaging* 10(3):620–629.
- P. Markiewicz, M. C. Goh (1994). Atomic force microscopy probe tip visualization and improvement of images using a simple deconvolution procedure. *Langmuir* 10:5–7.
- P. Markiewicz, M. C. Goh (1995). Atomic force microscope tip deconvolution using calibration arrays. *Rev. Sci. Instrum.* 66:3186–3190.
- D. Marr (1982). *Vision*. W. H. Freeman, San Francisco, CA.
- D. Marr, E. Hildreth (1980). Theory of edge detection. *Proc. R. Soc. Lond.* B207:187–217.
- D. Marr, T. Poggio (1976). Cooperative computation of stereo disparity. *Science* 194:283–287.
- G. A. Mastin (1985). Adaptive Filters for digital image noise smoothing: an evaluation. *Comput. Vis. Graph. Image Proc.* 31:102–121.
- A. D. McAulay, J. Wang, J. Li (1993). Optical wavelet transform classifier with positive real Fourier transform wavelets. *Optical Eng.* 32(6):1333–1339.
- W. S. McCulloch, W. Pitts (1943). A logical calculus of the ideas imminent in nervous activity. *Bull. Math. Biophys.* 5:115.
- J. J. Mecholsky, D. E. Passoja (1985). Fractals and Brittle Fracture, in *Fractal Aspects of Materials*. Materials Research Society, Pittsburgh PA.
- J. J. Mecholsky, T. J. Mackin, D. E. Passoja (1986). Crack propagation in brittle materials as a fractal process, in *Fractal Aspects of Materials II*. Materials Research Society, Pittsburgh, PA.
- J. J. Mecholsky, D. E. Passoja, K. S. Feinberg–Ringel (1989). Quantitative analysis of brittle fracture surfaces using fractal geometry. *J. Am. Ceram. Soc.* 72:60.
- G. Medioni, R. Nevatia (1985). Segment-based stereo matching. *Computer Vision, Graph. Image Process* 31:2–18.
- D. L. Milgram (1975). Computer Methods for Creating Photomosaics. *IEEE Trans.* C-24:1113–1119.
- D. L. Milgram, M. Herman (1979). Clustering edge values for threshold selection. *Comput. Graph. Image. Process* 10:272–280.
- M. Minsky, S. Papert (1969). *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA.
- M. W. Mitchell, D. A. Bonnell (1990). Quantitative topographic analysis of fractal surfaces by scanning tunneling microscopy. *J. Mater. Res.* 5(10):2244–2254.

- J. R. Monck, A. F. Oberhauser, T. J. Keating, et al. (1992). Thin-section ratiometric Ca²⁺ images obtained by optical sectioning of Fura-2 loaded mast cells. *J. Cell Biol.* 116:745-759.
- H. P. Moravec (1977). Towards automatic visual obstacle avoidance. *Proc. 5th IJCAI*, 584.
- R. B. Mott (1995). Position-tagged spectrometry, a new approach for EDS spectrum imaging. *Proc. Microsc. Microanal.* Jones & Begall, NY, 595.
- J. C. Mott-Smith (1970). Medial Axis Transformations, in *Picture Processing and Psychopictorics* (B. S. Lipkin, A. Rosenfeld, Eds.). Academic Press, New York.
- H. R. Myler, A. R. Weeks (1993). *Pocket Handbook of Image Processing Algorithms in C*. Prentice Hall, Englewood Cliffs, NJ.
- K. S. Nathan, J. C. Curlander (1990). Reducing speckle in one-look SAR images. *NASA Tech Briefs* Feb:70.
- F. Natterer (2001). *The Mathematics of Computerized Tomography*. SIAM, Philadelphia, PA.
- F. Natterer, F. Wubbeling (2001). *Mathematical Methods in Image Reconstruction*. SIAM, Philadelphia, PA.
- B. Neal, J. C. Russ, J. C. Russ (1998). A super-resolution approach to perimeter measurement. *J. Comput. Assisted Microsc.* 10(1):11-22.
- C. V. Negoita, D. A. Ralescu (1987). *Simulation, Knowledge-Based Computing, and Fuzzy Statistics*. Van Nostrand Reinhold, New York, NY.
- C. V. Negoita, D. A. Ralescu (1975). *Applications of Fuzzy Sets to Systems Analysis*. Halsted Press, New York, NY.
- R. Nevatia, K. Babu (1980). Linear feature extraction and description. *Comput. Graph. Image Process.* Vol. 13.
- W. Niblack (Ed.) (1993). Storage and Retrieval for Image and Video Databases. *SPIE Proc.* Vol. 1908.
- A. Nicoulin, M. Mattavelli, W. Li, M. Kunt (1993). Subband image coding using jointly localized filter banks and entropy coding based on vector quantization. *Optical Eng.* 32(7):1430-1450.
- A. Nieminen, P. Heinonen, Y. Nuevo (1987). A new class of detail-preserving filters for image processing. *IEEE Trans. PAMI* 9:74-90.
- N. Nikolaidis, I. Pitas (2001). *3-D Image Processing Algorithms*. John Wiley & Sons, New York.
- J. F. O'Callaghan (1974). Computing the Perceptual Boundaries of Dot Patterns. *Comp. Graph. Image Process.* 3(2):141-162.
- K. Oistämö, Y. Neuvo (1990). Vector median operations for color image processing. *Nonlinear Image Processing* (E. J. Delp, Ed.). *SPIE Proc.* 1247:2-12.
- C. K. Olsson (1993). Image Processing Methods in Materials Science, Ph.D. Thesis, Technical Univ. of Denmark, Lyngby Denmark.
- N. Otsu (1979). A threshold selection method from gray-level histograms. *IEEE Trans SMC* 9:62-69; 377-393.
- D. P. Panda, A. Rosenfeld (1978). Image segmentation by pixel classification in (gray level, edge value) space. *IEEE Trans. Comput.* 27:875-879.
- Y-H. Pao (1989). *Adaptive Pattern Recognition and Neural Networks*. Addison-Wesley, Reading, MA.
- J. R. Parker (1997). *Algorithms for Image Processing and Computer Vision*. John Wiley & Sons, New York.
- T. Pavlidis (1977). *Structural Pattern Recognition*. Springer Verlag, New York.
- T. Pavlidis (1980). A Thinning Algorithm for Discrete Binary Images. *Comput. Graph. Image Process.* 13:142-157.

- T. Pavlidis (1982). *Algorithms for Graphics and Image Processing*. Computer Science Press, Rockville, MD.
- L. D. Peachey, J. P. Heath (1989). Reconstruction from stereo and multiple electron microscope images of thick sections of embedded biological specimens using computer graphics methods. *J. Microsc.* 153:193–204.
- D. M. Pearsall (1978). Phytolith analysis of archaeological soils: evidence for maize cultivation in formative Ecuador. *Science* 199(4325):177–178.
- S. Peleg, J. Naor, R. Hartley, et al. (1984). Multiple resolution texture analysis and classification. *IEEE Trans. PAMI* 6:518.
- A. P. Pentland (1983). Fractal-based description of natural scenes *IEEE Trans PAMI*-6:661.
- A. P. Pentland, Ed. (1986). *From Pixels to Predicates*, Ablex, Norwood, NJ.
- A. P. Pentland, R. W. Picard, S. Sclaroff (1994). Photobook: Content-based manipulation of image databases, in *Proc. Storage Retrieval Image Video Databases II*, SPIE, San Jose, CA.
- E. Persoon, K.-S. Fu (1977). Shape Discrimination Using Fourier Descriptors. *IEEE Trans. Sys. Man Cyber. (SMC)* 7:170–179.
- J. L. Pfalz (1976). Surface networks. *Geographical Anal.* 8(2):77–93.
- D. R. Piperno (1984). A comparison and differentiation of phytoliths from maize (*Zea mays* L.) and wild grasses: use of morphological criteria. *American Antiquity* 49(2):361–383.
- I. Pitas (2000). *Digital Image Processing Algorithms and Applications*. John Wiley & Sons, New York.
- W. K. Pratt (1991). *Digital Image Processing*, 2nd edition. John Wiley & Sons, New York.
- T. Prettyman, R. Gardner, J. Russ, et al. (1991). On the performance of a combined transmission and scattering approach to industrial computed tomography, in *Advances in X-ray Analysis* Vol. 35. Plenum Press, New York, NY.
- J. M. S. Prewitt, M. L. Mendelsohn (1966). The analysis of cell images. *Ann. N.Y. Acad. Sci.* 128:1035–1053.
- L. Quam, M. J. Hannah (1974). *Stanford automated photogrammetry research AIM-254*. Stanford AI Lab, Stanford, CA.
- C. F. Quate (1994). The AFM as a tool for surface imaging. *Surf. Sci. (Netherlands)* 299–300, 980–95.
- J. Radon (1917). Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Berlin Sächsische Akad. Wissen.* 29:262–279.
- R. Ramanath (2000). Interpolation methods for the Bayer color array, M. S. Thesis, N. C. State Univ. Dept of Electrical Eng. (www4.ncsu.edu/~rramana/Research/MastersThesis.pdf).
- B. S. Reddy, B.N. Chatterji (1996). An FFT-based technique for translation, rotation, and Scale-Invariant image registration. *IEEE Trans. Image Process.* 5(8):1266–1271.
- M. G. Reed, C. V. Howard, C. G. Shelton (1997). Confocal imaging and second-order stereological analysis of a liquid foam. *J. Microsc.* 185(3):313–320.
- M. G. Reed, C. V. Howard (1997). Edge corrected estimates of the nearest neighbor function for three-dimensional point patterns. *J. Microsc.* (in press).
- M. G. Reed, C. V. Howard (1998). *Unbiased Stereology*. Bios Scientific Publ., Oxford.
- A. A. Reeves (1990). Optimized Fast Hartley Transform with Applications, in *Image Processing Thesis*, Dartmouth College, Hanover, NH.
- R. G. Reeves, Ed. (1975). *Manual of Remote Sensing*. American Society of Photogrammetry, Falls Church, CA.
- K. Rehm, S. C. Strother, J. R. Anderson, et al. (1994). Display of merged multimodality brain images using interleaved pixels with independent color scales. *J. Nucl. Med.* 35:1815–1821.

- I. Rezanaka, R. Eschbach, Eds. (1996). *Recent Progress in Ink-Jet Technologies*. Soc. Imaging Science and Tech., Springfield, VA.
- W. H. Richardson (1972). Bayesian-based iterative method of image restoration. *J. Opt. Soc. Am.* 62:55-59.
- J. P. Rigaut (1988). Automated image segmentation by mathematical morphology and fractal geometry. *J. Microsc.* 150:21-30.
- G. X. Ritter, J. N. Wilson (2001). *Handbook of Computer Vision Algorithms in Image Algebra*, 2nd edition. CRC Press, Boca Raton, FL.
- L. Roberts (1982). Recognition of three-dimensional objects, in *The Handbook of Artificial Intelligence*, Vol. III (P. Cohen, E. Figenbaum, Eds.). Kaufmann, Los Gatos, CA.
- L. G. Roberts (1965). Machine perception of three-dimensional solids, in *Optical and Electro-Optical Information Processing* (J. T. Tippett, Ed.) MIT Press, Cambridge, MA.
- F. J. Rohlf, J.W. Archie (1984). A comparison of Fourier methods for the description of wing shape in mosquitoes (Diptera culicidae). *Syst. Zool.* 33: 302-317.
- F. J. Rohlf (1990). Morphometrics. *Annu. Rev. Ecol. Syst.* 21: 299-316.
- D. W. Rolston (1988). *Principles of Artificial Intelligence and Expert System Development*. McGraw-Hill, New York.
- I. Rock (1984). *Perception*. W. H Freeman, New York.
- B. G. Rosen, R. J. Crafoord, Eds. (1997). *Metrology and Properties of Engineering Surfaces*. Chalmers Univ., Göteborg, Sweden.
- F. Rosenblatt (1958). The Perceptron: A probabilistic model for information organization and storage in the brain. *Psych. Rev.* 65:358-408.
- A. Rosenfeld, A. C. Kak (1982). *Digital Picture Processing*, Vols. 1 & 2. Academic Press, New York.
- I. Rovner (1971). Potential of opal phytoliths for use in paleoecological reconstruction. *Quaternary Res.* 1(#3):345-59.
- Y. Rui, T. S. Huang, S-F. Chang (1999). Image Retrieval: Current techniques, promising directions and open issues, *J. Visual Commun. Image Representation* 10(1):39-62.
- D. E. Rumelhart, G. E. Hinton, R. J. Williams (1986). Learning representations by back-propagating errors. *Nature* 323:533-536.
- J. C. Russ (1984). Implementing a new skeletonizing method. *J. Microsc.* 136:RP7.
- J. C. Russ (1986). *Practical Stereology*. Plenum Press, New York.
- J. C. Russ (1988). Differential Absorption Three-Dimensional Microtomography. *Trans Amer. Nucl. Soc.* 56(3):14.
- J. C. Russ (1990a). *Computer Assisted Microscopy*. Plenum Press, New York.
- J. C. Russ (1990b). Surface characterization: Fractal dimensions, Hurst Coefficients and Frequency transforms. *J. Comput. Assist. Microsc.* 2(3):161-184.
- J. C. Russ (1990c). Processing images with a local Hurst operator to reveal textural differences. *J. Comput. Assist. Microsc.* 2(4):249-257.
- J. C. Russ (1991). Multiband thresholding of images. *J. Comput. Assist. Microsc.* 3(2):77-96.
- J. C. Russ (1993). JPEG Image Compression and Image Analysis. *J. Comput. Assist. Microsc.* 5(3):237-244.
- J. C. Russ (1993). Method and application for ANDing features in binary images. *J. Comput. Assist. Microsc.*, 5(4):265-272.
- J. C. Russ (1994). *Fractal Surfaces*, Plenum Press, New York.
- J. C. Russ (1995a). Computer-Assisted Manual Stereology. *J. Comp. Assist. Microsc.* 7(1):35-46.
- J. C. Russ (1995b). Median Filtering in Color Space. *J. Comp. Assist. Microsc.* 7(2):83-90.
- J. C. Russ (1995c). Thresholding Images. *J. Comp. Assist. Microsc.* 7(3):41-164.

- J. C. Russ (1995d). Designing kernels for image filtering. *J. Comp. Assist. Microsc.* 7(4):179-190.
- J. C. Russ (1995e). Optimal greyscale images. *J. Comp. Assist. Microsc.* 7(4):221-234.
- J. C. Russ (1995f). Segmenting touching hollow features. *J. Comp. Assist. Microsc.* 7(4):253-261.
- J. C. Russ (1997). Fractal dimension measurement of engineering surfaces, in *7th Intl Conf. on Metrology Properties Eng. Surf.* (B. G. Rosen, R. J. Crafoord, Eds.). Chalmers Univ., Göteborg Sweden, 170-174.
- J. C. Russ (2001a). *Forensic Uses of Digital Microscopy*, CRC Press, Boca Raton, FL.
- J. C. Russ (2001b). Fractal Geometry in Engineering Metrology, in *Metrology and Properties of Engineering Surfaces* (E. Mainsah et al., Eds.). Kluwer Academic Publishers, London, 43-82.
- J. C. Russ, D. S. Bright, J. C. Russ, et al. (1989). Application of the Hough transform to electron diffraction patterns. *J. Comput. Assist. Microsc.* 1(1):3-77.
- J. C. Russ, R. T. Dehoff (2001). *Practical Stereology*, 2nd Edition. Plenum Press, New York.
- J. C. Russ, I. Rovner (1987). Stereological verification of Zea Phytolith taxonomy. *Phytolitharien Newsletter* 4(#3):10.
- J. C. Russ, J. C. Russ (1988a). Automatic discrimination of features in grey scale images. *J. Microsc.* 148:263-277.
- J. C. Russ, J. C. Russ (1988b). Improved implementation of a convex segmentation algorithm. *Acta Stereologica* 7:33-40.
- J. C. Russ, J. C. Russ (1989a). Uses of the Euclidean Distance Map for the Measurement of Features in Images. *J. Comput. Assist. Microsc.* 1(4):343.
- J. C. Russ, J. C. Russ (1989b). Topological Measurements on Skeletonized Three-Dimensional Networks. *J. Comput. Assist. Microsc.* 1:131-150.
- J. C. Russ, H. Palmour, III, T. M. Hare (1989). Direct 3-D Pore Location Measurement in Alumina. *J. Microsc.* 155(2):RP1-2.
- P. Russell, D. Batchelor (2001). SEM and AFM: Complementary techniques for surface investigations. *Microsc. Anal.* July 2001: 5-8.
- F. F. Sabins, Jr. (1987). *Remote Sensing: Principles and Interpretation*, 2nd Edition, W. H. Freeman, New York.
- K. Sahoo, S. Soltani, A. C. Wong, et al. (1988). A survey of thresholding techniques. *Comput. Vision Graph. Image Process.* 41:233-260.
- E. Sanchez, L. A. Zadeh, Eds. (1987). *Approximate Reasoning in Intelligent System Decision and Control*. Oxford Press, New York NY.
- J. Sanchez, M. P. Canton (1999). *Space Image Processing*. CRC Press, Boca Raton, FL.
- L. J. Sartor, A. R. Weeks (2001). Morphological operations on color images. *J. Elect. Imaging* 10(2):548-559.
- S. Saunders (1991). Magellan: the geologic exploration of venus. *Eng. Sci.* Spring 1991:15-27.
- A. Savitsky, M. J. E. Golay (1964). Smoothing and differentiation of data by simplified least squares procedures. *Anal. Chem.* 36:1627-1639.
- R. J. Schalkoff (1991). *Pattern Recognition: Statistical, Syntactical and Neural Approaches*. John Wiley & Sons, New York.
- D. J. Schnerberk, H. E. Martz, S. G. Azavedo (1991). Multiple Energy Techniques in Industrial Computerized Tomography, in *Review of Progress in Quantitative Nondestructive Evaluation* (D. O. Thompson, D. E. Chimenti, Eds.), Plenum Press, New York.
- H. P. Schwartz, K. C. Shane (1969). Measurement of particle shape by Fourier analysis. *Sedimentology* 13:213-231.

- H. Schwarz, H. E. Exner (1980). Implementation of the concept of fractal dimensions on a semi-automatic image analyzer. *Powder Techn.* 27:107.
- H. Schwarz, H. E. Exner (1983). The characterization of the arrangement of feature centroids in planes and volumes. *J. Microsc.* 129:155.
- P. J. Scott (1995). Recent advances in areal characterization. *IX Intern. Oberflächenkolloq.* Tech. Univ. Chemnitz-Zwickau, 151–158.
- P. J. Scott (1997). Foundations of topological characterization of surface texture in *7th Intl Conf. on Metrology and Properties of Engineering Surfaces* (B. G. Rosen, R. J. Craford, Eds.), Chalmers Univ., Göteborg, Sweden, 162–169.
- J. Serra (1982). *Image Analysis and Mathematical Morphology*. Academic Press, London.
- M. Seul, L. O’Gorman, M. J. Sammon (2000). *Practical Algorithms for Image Analysis*. Cambridge Univ. Press, Cambridge, MA.
- L. A. Shepp, B. F. Logan (1974). The Fourier reconstruction of a head section. *IEEE Trans. NS* 21:21–43.
- A. Shih, G. Wang, P. C. Cheng (2001). Fast algorithm for X-ray cone-beam microtomography. *Microsc. Microanal.* 7:13–23.
- J. Skilling (1986). Theory of Maximum Entropy Image Reconstruction, in *Maximum Entropy and Bayesian Methods in Applied Statistics, Proc. 4th Max Entropy Workshop, Univ. of Calgary, 1984* (J. H. Justice, Ed.). Cambridge Univ. Press, Cambridge, MA. 156–178.
- P. E. Slatter (1987). *Building Expert Systems: Cognitive Emulation*. Halsted Press, New York.
- B. D. Smith (1990). Cone-beam tomography: recent advances and a tutorial review. *SPIE Optical Eng.* 29:5.
- R. F. Smith (1990). *Microscopy and Photomicrography*, CRC Press, Boca Raton, FL.
- B. Smolka, A. Chydzinski, K. W. Wojciechowski, et al. (2001). On the reduction of impulse noise in multichannel image processing. *Opt. Eng.* 40(6):902–908.
- D. L. Snyder, T. J. Schutz, J. A. O’Sullivan (1992). Deblurring subject to nonnegative constraints. *IEEE Trans. Sig. Proc.* 40:1143–1150.
- I. Sobel (1970). *Camera models and machine perception AIM-21*. Stanford Artificial Intelligence Lab, Palo Alto, CA.
- P. Soille (1999). *Morphological Image Analysis*, Springer-Verlag, Berlin.
- M. Sonka, V. Hlavac, R. Boyle (1999). *Image Processing, Analysis and Machine Vision*, 2nd edition, Brooks Cole, Pacific Grove, CA.
- S. Srinivasan, J. C. Russ, R. O. Scattergood (1990). Fractal Analysis of Erosion Surfaces. *J. Mater. Res.* 5(11):2616–2619.
- J. A. Stark, W. J. Fitzgerald (1996). An Alternative Algorithm for Adaptive Histogram Equalization. *Comput. Vis. Graph. Image Proc.* 56(2):180–185.
- M. Stefik (1995). *Introduction to Knowledge Systems*. Morgan Kaufmann, San Francisco, CA.
- D. C. Sterio (1984). The unbiased estimation of number and sizes of arbitrary particles using the disector. *J. Microsc.* 134:127–136.
- P. L. Stewart, R. M. Burnett Seminars in Virology, cited in C. J. Mathias, Visualization Techniques Augment Lab Research into Structure of Adenovirus. *Scientific Computing Automation* 7(6):51–56.
- M. C. Stone, W. B. Cowan, J. C. Beatty (1988). Color gamut mapping and the printing of digital color images, *ACM Trans. Graphics* 7(4):249–292.
- J. A. Storer (1992). *Image and Text Compression*. Kluwer Academic Publishers, New York.
- K. J. Stout, P. J. Sullivan, W. P. Dong, et al. (1993). *The development of methods for the characterization of roughness in three dimensions*. Publication EUR 15178 EN of the Commission of the European Communities, Univ. of Birmingham, Edgbaston, England.

- R. G. Summers, C. E. Musial, P-C. Cheng, et al. (1991). The use of confocal microscopy and stereocon reconstructions in the analysis of sea urchin emryonic cell division. *J. Electron Microscope Tech.* 18:24–30.
- R. E. Swing (1997). *An Introduction to Microdensitometry*. SPIE Press, Bellingham, WA.
- H. Talbot, T. Lee, D. Jeulin, et al. (2000). Image analysis of insulation mineral fibers. *J. Microsc.* 200(3):251–258.
- T. R. Thomas (1999). *Rough Surfaces*, 2nd edition. Imperial College Press, London.
- M. M. Thompson, Ed. (1966). *Manual of Photogrammetry*, American Society of Photogrammetry, Falls Church, VA.
- J. T. Tou, R. C. Gonzalez (1981). *Pattern Recognition Principles*. Addison-Wesley, Reading, MA.
- J. Trussell (1979). Comments on "Picture thresholding using an iterative selection method." *IEEE Trans SMC* 9:311.
- T-M. Tu, S-C. Su, H-C. Shyu, et al. (2001). Efficient intensity-hue-saturation-based image fusion with saturation compensation. *Optical Eng.* 40(5):720–728.
- J. N. Turner, D. H. Szaeowski, K. L. Smith, et al. (1991). Confocal microscopy and three-dimensional reconstruction of electrophysiologically identified neurons in thick brain slices. *J. Electron Microscope Tech.* 8:11–23.
- P. C. Twiss, E. Suess and R.M. Smith (1969). Morphological classification of grass phytoliths. *Soil Science Soc. America Proc.* 33(1):109–115.
- S. E. Umbaugh (1998). *Computer Vision and Image Processing*. Prentice Hall, Saddle River NJ.
- E. E. Underwood (1970). *Quantitative Stereology* Addison-Wesley, Reading, MA.
- E. E. Underwood, K. Banerji (1986). Fractals in Fractography. *Mater. Science Eng.* 80:1.
- P. A. van den Elsen, J. B. A. Maintz, E-J.D. Pol, et al. (1995). Automatic Registration of CT and MR Brain Images Using Correlation of Geometrical Features. *IEEE Trans. Med. Imaging*, 14(2):384–396.
- P. A. van den Elsen, E-J. D. Pol, T. S. Sumanaweera, et al. (1994). Grey value correlation techniques used for automatic matching of CT and MR brain and spine images. *SPIE Intl Conf. Visualization in Biomedical Computing* 2359:227–237.
- P. A. van den Elsen, E-J. D. Pol, M.A. Viergever (1993). Medical image matching — a review with classication, *IEEE Eng. Medicine Biol.* 12(1):26–39.
- J. Van Helden (1994). CrestPathway algorithm. (personal communication).
- H. R. J. R. Van Helleputte, T. B. J. Haddeman, M. J. Verheijen, J-J. Baalbergen (1995). Comparative study of 3D measurement techniques (SPM, SEM, TEM) for submicron structures. *Microelectron. Eng.* 27:547.
- G. M. P. Van Kempen, L. J. Van Vliet, P. J. Verveer, et al. (1997). A quantitative comparison of image restoration methods in confocal microscopy. *J. Microsc.* 185(3):354–365.
- J. G. Verly, R. L. Delanoy (1993). Some principles and applications of adaptive mathematical morphology for range imagery. *Optical Eng.* 32(12):3295–3306.
- H. Verschueren, B. Houben, J. De Braekeleer, et al. (1993). Methods for computer assisted analysis of lymphoid cell shape and motility, including Fourier analysis of cell outlines. *J. Immunol. Meth.* 163:99–113.
- J. S. Villarrubia (1994). Morphological estimation of tip geometry for scanner probe microscopy. *Surf. Sci.* 321:287–300.
- J. S. Villarrubia (1996). Scanned probe microscope tip characterization without calibrated tip characterizers. *J. Vac. Sci. Technol.* B14:1518–1521.
- J. S. Villarrubia (1997). Algorithms for scanned probe microscopy image simulation, surface reconstruction and tip estimation. *J. Res. Natl. Inst. Stand. Tech.* 102(4):425.

- R. J. Wall, A. Klinger, K. R. Castleman (1974). Analysis of Image Histograms. *Proc. 2nd Joint Intl Conf. Patt. Recog. IEEE 74CH-0885-4C*, 341–344.
- J. R. Walters (1988). *Crafting Knowledge-Based Systems: Expert Systems Made Easy*. John Wiley & Sons, New York.
- G. Wang, T. H. Lin, P. C. Cheng, et al. (1991). Scanning cone-beam reconstruction algorithms for X-ray microtomography. *SPIE Scanning Microscope Instrumentation* 1556:99.
- G. Wang, M. Vannier (2001). Micro-CT scanners for biomedical applications, *Advanced Imaging* July:18–27.
- Z. Wang (1990). *Principles of Photogrammetry (with Remote Sensing)*. Press of Wuhan Technical University of Surveying and Mapping, Beijing.
- A. R. Weeks (1996). *Fundamentals of Electronic Image Processing*, SPIE Press, Bellingham, WA.
- E. R. Weibel (1979). *Stereological Methods*, Vols. I & II. Academic Press, London.
- S. Welstead (1999). *Fractal and Wavelet Image Compression Techniques*. SPIE Press, Bellingham, WA.
- A. Wen, C. Lu (1993). Hybrid vector quantization *Optical Eng.* 32(7):1496–1502.
- J. West, J. M. Fitzpatrick, M.Y. Wang, et al. (1997). Comparison and Evaluation of Retrospective Intermodality Registration Techniques. *J. Comput. Assist. Tomography* 21:54–566.
- J. S. Weszka (1978). A survey of threshold selection techniques. *Comput. Graph. Image. Process* 7:259–265.
- J. Weszka, C. Dyer, A. Rosenfeld (1976). A comparative study of texture measures for terrain classification. *IEEE Trans SMC* 6:269–285.
- J. S. Weszka, A. Rosenfeld (1979). Histogram modification for threshold selection. *IEEE Trans SMC* 9:38–52.
- D. J. Whitehouse (1994). *Precision — the Handbook of Surface Metrology*. Institute of Physics Publishing, Bristol.
- H. K. Wickramasinghe (1991). Scanned probes old and new. *AIP Conf. Proc.* 9–22.
- B. Willis, B. Roysam, J. N. Turner, et al. (1993). Iterative, constrained 3D image reconstruction of transmitted light bright field micrographs based on maximum likelihood estimation, *J. Microsc.* 169:347–361.
- R. Wilson, M. Spann (1988). *Image Segmentation and Uncertainty*. John Wiley & Sons, New York.
- G. Winstanley (1987). *Program Design for Knowledge-Based Systems*, Halsted Press, New York.
- R. J. Woodham (1978). Photometric stereo: a reflectance map technique for determining surface orientation from image intensity. *Proc. SPIE* 155: 136–143.
- G. Wolf (1991). Usage of global information and a priori knowledge for object isolation. *Proc. 8th Intl. Congr. for Stereology*, Irvine, CA., 56.
- K. W. Wong (1980). Basic Mathematics of Photogrammetry, in *Manual of Photogrammetry*, 4th Edition. American Society of Photogrammetry, Falls Church, VA., 57–58.
- B. P. Wrobel (1991). Least-squares methods for surface reconstruction from images. *ISPRS J. Photogrammetry Remote Sensing*, 46:67–84.
- S. Wu, A. Gersho (1993). Lapped vector quantization of images. *Optical Eng.* 32(7):1489–1495.
- R. R. Yager (1979). On the measures of fuzziness and negation, part 1: membership in the unit interval. *Int. J. Gen. Sys.* 5:221–229.
- G. J. Yang, T. S. Huang (1981). The effect of median filtering on edge location estimation. *Comput. Graph. Image Proc.* 15:224–245.

- N. Yang, J. Boselli, I. Sinclair (2001). Simulation and quantitative assessment of homogeneous and inhomogeneous particle distributions in particulate metal matrix composites. *J. Microsc.* 201(2):189–200.
- Y. Yakimovsky (1976). Boundary and object detection in real world images. *J. Assoc. Comput. Mach.* 23:599–618.
- R. W. Young, N. G. Kingsbury (1993). Video compression using lapped transforms for motion estimation/compensation and coding. *Optical Eng.* 32(7):1451–1463.
- T. York (2001). Status of electrical tomography in industrial applications. *J. Electron. Imaging* 10(3):608–619.
- A. Yoshitaka, T. Ichikawa (1999). A survey on content-based retrieval for multimedia databases. *IEEE Trans. Knowledge Data Eng.* 11(1):81–93.
- L. A. Zadeh (1965). Fuzzy Sets. *Infor. Control* 8:338–353.
- C. T. Zahn, R.Z. Roskies (1972). Fourier descriptors for plane closed curves. *IEEE Trans. Comput.* C-21:269–281.
- X. Zhou, E. Dorrer (1994). An Automatic Image Matching Algorithm Based on Wavelet Decomposition. *ISPRS Int. Arch. of Photogrammetry Remote Sensing*, 30(3/2):951–960.
- H-J. Zimmermann (1987). *Fuzzy Sets, Decision Making and Expert Systems*. Kluwer Academic Publishers, Boston, MA.
- S. Zucker (1976). Region growing: childhood and adolescence. *Comput. Graph. Image Process* 5:382–399.
- O. Zuniga, R. M. Haralick (1983). Corner detection using the facet model, *Comput. Vision Patt. Recog. IEEE*, 30–37.

Index

A

- Abbasi-Dezfouli and Freeman studies, 660
- Abbott curve, 685
- Absolute colors, 20, 45, *see also* Colors
- Absorption tomography, 558
- Accumulator spaces, *see* Hough spaces
- Accuracy, 357–359
- Acoustic microscope images, 2, 628
- Acquiring images, *see also* Imaging
 - bandwidth limitations, 18–20
 - cameras, 8–18, 36–39
 - color images, 30–52
 - electronics, 18–20
 - focusing, 18
 - grey-scale resolution, 23–24
 - high-depth images, 26–29
 - image types, 52–54
 - imaging requirements, 71–77
 - limitations, 36–39
 - multiple images, 59–64
 - noise, 24–26
 - pixels, 20–23
 - range imaging, 54–59
 - stereoscopy, 64–71
 - video cameras, 6–7
- Active contours, 369
- Adaptive equalization, 219
- Additivity, 285
- Aerial photography
 - color shading, 181
 - cross-correlation method, 110–111
 - feature-matching, 502
 - geographical information system (GIS), 108
 - geometrical distortion, 192
 - remote sensing classification, 534
 - stereoscopy, 66–69
 - subtraction, 268
- AFIS, *see* Automated Fingerprint Identification System (AFIS)
- AFMs, *see* Atomic force microscopes (AFMs)
- Agfa studies, 90
- Aguilar studies, 58
- Aigeltinger studies, 649
- Airplane photography, *see* Aerial photography
- Algebraic reconstruction technique (ART), 563–567, 583
- “Alice’s Restaurant,” 79
- Alignment, 193–197
- Alpha channel, 389
- American Society for Testing and Materials (ASTM), 453–456
- Amplifier noise reduction, 13
- Amplitude parameters, 687
- Anderberg studies, 669
- Animal vision, 1, 64–65
- Anisotropy, 416, 466–469
- Arce studies, 158
- Archival properties, 96, 103, 104
- Area fraction, 447
- ART, *see* Algebraic reconstruction technique (ART)
- Artworks, 186
- Arvo studies, 506
- Astigmatism, 292–293
- ASTM, *see* American Society for Testing and Materials (ASTM)
- Astola, Dougherty and, studies, 409
- Astolo studies, 154
- Astronomical imaging, *see also* Aerial photography; Satellite imaging; Weather imaging
 - alignment, 195
 - blinking, 268
 - charge-coupled device (CCD) cameras, 13
 - Clementine orbiter, 22
 - color substitution, 32
 - compression, 304
 - contrast range, 27
 - focus, 318
 - galaxy M51, 242
 - high-resolution imaging, 589
 - Hubble telescope, 10, 318
 - image dependence, 1
 - Laplacian methods, 224–225
 - Magellan probe, 69–70, 632
 - planets and asteroids identification, 268
 - Saturn, 224, 304

- Sobel operator, 242
- Venus, 632
- Atomic force microscopes (AFMs)
 - contours, 371
 - distance recording, 54
 - image types, 54
 - interferometer comparison, 58–59
 - noncontacting measurements, 657
 - physical contact, 654
 - quantitative dimensional measurements, 656
 - range imaging, 55–58, 665–668
 - rank operations, 255
 - surface data analysis, 681
- Atom probe ion microscope, 600
- Autocorrelation process, 331–332, 685, 690–691, *see also* Correlation process; Cross-correlation process
- Automated Fingerprint Identification System (AFIS), 109
- Automatic scaling, 265

B

- Baba studies, 591
- Babu, Nevatia and, studies, 435
- Background fitting, 174–179
- Background subtraction, 173, *see also* Subtraction
- Backprojection, 560–561, 582–584
- Backup copy storage, 104
- Baddeley studies, 452, 487
- Baker, Hearn and, studies, 190, 606
- Balasubramanian studies, 117
- Ballard and Brown studies, 376, 529
- Ballard studies, 491, 529
- Balloons, 379
- Bandwidth limitations, 18–20
- Banerji, Underwood and, studies, 261
- Barnard and Thompson studies, 660
- Barnes studies, 570
- Barnsley and Hurd studies, 126, 127
- Barnsley and Sloan studies, 127
- Barnsley studies, 127
- Barty studies, 596
- Batchelor, Russell and, studies, 654
- Bayer pattern, 15
- Bayes' theorem, 168, 323, 536, 567
- BCC, *see* Body-centered cubic (BCC) configuration
- Beam hardening, 573–578
- Bertrand studies, 253
- Berzins studies, 229
- Besl studies, 606
- Beucher, Lantjeoul and, studies, 430
- Beucher and Lantjeoul studies, 430
- Bezier curves, 508
- Bially iteration, 323
- Biamond, Legendijk and, studies, 324
- Binnig studies, 654
- Binning, 11
- Biomedical imaging, 535–536
- Birmingham measurement suite, 687–693
- Blackie studies, 623
- Black studies, 549
- "Blinking" images, *see* Astronomical imaging
- Blue channel, 16
- Blurring, *see also* Smoothing
 - deconvolution, 319, 322–324
 - frequency space, 324–327
 - horizontal line readout, 13
 - motion, 324–327
 - neighborhood averaging, 142
 - reconstruction, 561–562
 - uniformity, 321
- Boddeke studies, 18, 476
- Body-centered cubic (BCC) configuration, 605
- Bonnell, Mitchell and, studies, 296
- Boolean logic
 - image math, 264
 - images (binary), 264, 383–389
 - measurement templates, 396
 - multiple images, 59
 - searching databases, 105, 109
 - textural orientation, 353
 - thresholding, 337, 344, 349, 352
 - visualization software programs, 346
- Borgefors studies, 647
- Boundaries, *see also* Edges
 - asymmetric, 359
 - images (binary), processing, 439–442
 - line representation, 358, 367–370
 - pixels, 22–23
 - power spectrum, 295
 - taut-string (rubber-band), 504
- Bounding box, 481
- Bow studies, 547
- Boyde studies, 66, 68
- Bracegirdle, Bradbury and, studies, 52
- Bracegirdle studies, 52
- Bracewell studies, 278
- Bradbury and Bracegirdle studies, 52
- Bradbury studies, 52
- Braudaway studies, 117
- Breadth, *see* Caliper dimensions
- Bresenham line drawing algorithm, 616
- Briarty and Jenkins studies, 591
- Bright and Steel studies, 250
- Brightness
 - color scales substitution, 31
 - contrast expansion, 132–133
 - histogram equalization, 214
 - human perception, 4, 9, 23
 - imaging requirements, 73
 - large range, cameras, 26
 - measurements, feature specific, 475–480
 - pixels, 21
 - solid-state chip cameras, 11
 - surface curvature, 176
- Bright studies, 669
- Broadcast television, 18–19, 129, 205
- Brodatz studies, 533
- Brown, Ballard and, studies, 376, 529
- Brown studies, 606
- Browsing, 108–111
- Brute-force correlation methods, 65
- Bryan and Skilling studies, 170
- Buckytubes, 665
- Burnett, Stewart and, studies, 585

Butterworth filters, 299–301, 306
Buzuloiu studies, 219

C

- CAD, *see* Computer graphics
Calibration scales, 478, 481
Caliper dimensions, 507–509
Cameras, *see also* Video cameras
 artefacts and limitations, 12–13
 charge-coupled device (CCD) cameras, 8–11
 color, 13–15, 35
 complementary metal oxide semiconductor (CMOS) cameras, 17–18
 digital, limitations, 18
 focusing, 18
 large brightness range, 26
 linearity, 476
 noisy images, 138
 performance level perspective, 27
 resolution, 15–16
 solid-state chip, 8–11
 studio, 11
 tube-type cameras, 7
Camouflage effectiveness, 5
Canny operators, 247–248, 510
Canton, Sanchez and, studies, 73
Capowski, Johnson and, studies, 591
Carasso studies, 324
Carlsen studies, 600
Carrington studies, 324
Castle and Zhdan studies, 654
Castleman studies, 380, 531
Cathode ray tube (CRT) monitors, 50–52, 88
CCD cameras, *see* Charge-coupled device (CCD) cameras
CDs, *see* Compact disks (CDs)
Cederberg studies, 374
Censor studies, 564
Centroids, 482–484, 488
Cesar, Costa and, studies, 73
Chain code
 harmonic analysis, 520–521
 perimeter, 509–510
 three-dimensional measurements, 524
 topology, 523
Change tree, surfaces, 693–694
Charge-coupled device (CCD) cameras
 basics, 8–11
 compression, 121
 high-resolution imaging, 589
 noisy images, 138
Chatterji, Reddy and, studies, 606
Cheerios, 556
Chermant, Coster and, studies, 409, 420
Chieco studies, 476
Chip cameras, *see* Charge-coupled device (CCD) cameras
Chord encoding (run-length encoding), 110, 116, 373–374, 392
Christus, Petrus, 186
Chui studies, 123
CIE, *see* Commission Internationale de L'Éclairage (CIE) standards
Circle and ellipse measurements, 505–506
Van Cittert iterative technique, 323
Classical erosion, *see* Erosion
Classification, *see also* Feature recognition
 fuzzy logic, 107
 imaging requirements, 76
 parameters, 531–534
 reliability, 532
 segmentation and thresholding, 379–381
Cleanliness, surface, 653
Clementine orbiter, 22
Closing (dilation/erosion), 410–413
Clusters, 379–380, 487, 545–548
CMOS cameras, *see* Complementary metal oxide semiconductor (CMOS) cameras
CMY/CMYK color
 color printing, 89–93
 image analysis, 45
 printing hardware, 93, 96
Cockett studies, 550
CODECs, *see* Coding/decoding (codec); Compressor/DECompressors (CODECs)
Coding/decoding (codec), 113–116, *see also* Compression
Cohen studies, 647
Colocalization, 383–384
Color cameras, *see* Cameras
Colorimetry, 185
Color printing, 88–93, 94
Colors, *see also* Absolute colors
 basics, 30–36
 converting to grey-scale, 46
 correction, 47–50, 185
 displays, 50–52
 human perception, 4
 imaging requirements, 73
 printing, 88–93
 pseudo-color, 31–32
 ray tracing, 625
 reduced palettes, 117–118
 separations, 91
 shading, 181–187
 subsampling, 33
 three-dimensional images, 615
 video camera limitations, 20
Color spaces, 40–47
Comer and Delp studies, 154
Comet tails, 321
Commission Internationale de L'Éclairage (CIE) standards
 color printing, 88–89
 color space, 41
 two-dimensional thresholds, 340
Compact disks (CDs), 101–103, 129
Complementary metal oxide semiconductor (CMOS) cameras, 17–18
Composition imaging, 662–664
Composition maps, 668–669
Compression
 delta, 112
 digital camera limitations, 36–37
 fractal, 126–128, 130
 frequency transform, 304

- JPEG, 118–123
 - lossless coding, 111–117
 - suitable, 126
 - three-dimensional image storage, 605
 - thumbnails, 108–111
 - wavelet, 123–126
 - Compressor/DECompressors (CODECs), 128
 - Computer graphics
 - defects, correcting, 189–191
 - non-planar views, 188
 - Phong shading, 673, 676
 - surfaces, 633–634, 636
 - Computerized tomography (CT) scanners, 460, 462, 580
 - Computer monitors, 15–16, 50–52
 - Condeco studies, 197
 - Cone-beam geometry, 587–589
 - Confocal scanning light microscope (CSLM)
 - color coding effects, 628
 - contours, 371
 - data sets, 610
 - deconvolution, 324
 - image math, 271
 - optical sectioning, 596–598
 - range imaging, 54
 - reflection, 628, 630
 - stereo viewing, 619
 - surfaces, 635
 - technology comparison, 28–29
 - volumetric imaging, 462
 - Contacting measurements, 653–657
 - Contour maps, 671
 - Contours, 370–373
 - Contrast, 131–136, 208–211
 - Convex segmentation, 500
 - Convex surfaces, 393
 - Convolution process
 - derivative operations, 244
 - frequency space, 314–332
 - rank operations, 251
 - Cookson studies, 624
 - Cooley and Tukey studies, 278
 - Cormack studies, 559
 - Cornsweet studies, 225
 - Correlation process, 314–332, 328–331, *see also*
 - Autocorrelation process; Cross-correlation process
 - Costa and Cesar studies, 73
 - Coster and Chermant studies, 409, 420
 - Counting features, 496–502
 - Covering blanket method, 696, 698
 - Crafoord, Rosen and, studies, 685
 - Cray supercomputer reconstructions, 589
 - Cross-correlation process, *see also* Autocorrelation process; Correlation process
 - data matching approach, 110
 - military targets, 529
 - shape and orientation, 252
 - template matching, 527–529
 - Cruz-Orive studies, 471
 - CSLM, *see* Confocal scanning light microscope (CSLM)
 - CT scanners, *see* Computerized tomography (CT) scanners
 - Cuneiform tablet shading, 674
 - Curds, 259
 - Curlander, Nathan and, studies, 165
 - Curvature, 176, 464–466, *see also* specific types of curves
 - Custer, 425
 - Cycloids, 453, 458, 467
- ## D
- "Daisy wheel" printers, 80
 - Van Dam, Foley and, studies, 190, 606
 - Danielsson studies, 427
 - Darkroom techniques, 27
 - Data analysis of surfaces, 677–683
 - Databases, 104–108
 - Data sets, 26, 604–611
 - Data sources, 591–592
 - Daubechies studies, 123
 - Daut studies, 112
 - Davidson studies, 435
 - Davies studies, 157
 - Davis, Freeman and, studies, 521
 - Davoust, DeMandolx and, studies, 669
 - Debye-Scherrer x-ray camera, 267, 478
 - Decision points, 534–536
 - Deckman studies, 587
 - Deconvolution, 319, 322–324
 - Defects
 - alignment, 193–197
 - background fittings, 174–179
 - color shading, 181–187
 - computer graphics, 189–191
 - contrast expansion, 131–136
 - geometrical distortion, 191–193
 - interpolation, 167–168, 197–202
 - maximum entropy methods, 168–170
 - maximum likelihood methods, 170–171
 - morphing, 202–205
 - neighborhood techniques, 140–167
 - noise, 137–140
 - non-planar views, 188–189
 - nonuniform illumination, 171–174
 - rank leveling, 179–181
 - three-dimensional images, 568–572
 - Dehoff, Russ and, studies, 445
 - Dehoff and Rhines studies, 445
 - Delanoy, Verly and, studies, 257
 - Delp, Comer and, studies, 154
 - Delta compression, 112
 - DeMandolx and Davoust studies, 669
 - Denley studies, 58
 - Deposition, surface build up, 652–653
 - Derivative operations
 - convolution, 314
 - image enhancement, 229–232
 - rank operations, 249
 - Determining numbers (measurements), 461–464
 - Dietzsch studies, 685
 - Difference of Gaussians (DOG), 243, 247
 - Digital cameras, 36–39, *see also* Cameras
 - Digital movies, 128–130, 160, 205
 - Digital video disks (DVDs), 102
 - Digital video recorders, 36

Dilation

- basics, 409–413
 - Boolean operations, 425
 - measurements, 417–418
 - neighborhoods, 420
 - noise reduction, 166
 - position information, 361
 - rank operations, 255
- Dinstein, Haralick and, studies, 336–337
- Discrete cosine transform (DCT), 118, 125–126
- Display hardware, 623–624
- Distance fade, 627
- Distortion, 191–193, 199
- Dithering, 87–88
- Divijver and Kittler studies, 527
- Division, 270–276, *see also* Segmentation
- DOG, *see* Difference of Gaussians (DOG)
- Doppler shift, 70
- Dorrer, Zhou and, studies, 660
- Dots in printing, 80, 83–88
- Double thresholding, 405–408
- Dougherty and Astola studies, 409
- Dry ink printers, 95
- Duda and Hart studies, 491, 527
- Dust, 72
- Dye-sublimation (diffusion) printers, 94–95
- Dynamic ranging, 265

E

- Edges, *see also* Boundaries
- enhancing operator, 242
 - following, 368–369, 379
 - segmentation and thresholding, 379
 - Sobel and Kirsch operators, 232–248
- Edwards studies, 149
- Ehrlich studies, 417
- Eight-connectedness, 376, 436
- Electron microscopy, 192, 288, *see also* Scanning
- electron microscopes (SEMs); Transmission
 - electron microscopes (TEMs) and imaging
- Electrons
- diffraction patterns, 2
 - noise from transfers, 13
 - shifting, 8
 - well capacity, 11
- Elevation, 32, 59, 69
- Van den Elsen studies, 606
- Embossed appearance, 672
- Emerging methods, 693–694
- Emission tomography, 577, *see also* Positron-emission tomography (PET)
- Engel and Massalski studies, 584
- Entropy methods, 115, 362, *see also* specific type of entropy method
- Erosion
- basics, 409–413
 - Boolean operations, 425
 - measurements, 417–418
 - neighborhoods, 420
 - noise reduction, 166
 - position information, 361
- Errors (dust related), 72, *see also* Defects
- Euclidean distance map (EDM)

- caliper dimensions, 509
- edge-to-edge distances, 489
- images (binary), processing, 425–444
- shape, 515
- standardized Euclidean distance, 539

- Exner, Schwarz and, studies, 486
- Expert systems, 549–550
- Extremum filter, 158

F

- Face-centered cubic (FCC) configuration, 605
- Fahmy studies, 261
- Faloutsos studies, 106
- Fan-beam geometry, 578, 579
- Fast Fourier transform (FFT), 278, 300, *see also* Fourier transforms
- Feature recognition, *see also* Classification
- cross-correlation, 527–529
 - decision points, 534–536
 - expert systems, 549–550
 - images (binary), processing, 391–405
 - kNN and cluster analysis, 545–548
 - learning systems, 541–545
 - multidimensional classification, 536–541
 - neural nets, 551–553
 - parametric description, 529–534
 - Sobel operator, 241
 - syntactical models, 553–554
 - template matching, 527–529
- Feature-specific measurements, *see* Measurements, feature specific
- Feder studies, 261, 296, 417
- Feldkamp studies, 583
- Feret's diameter, 484
- Fern image, 126–127
- FFT, *see* Fast Fourier transform (FFT)
- Files, 74, 100–101
- Film, 9, 11, 16
- Film recorders, 98–100
- Filters, *see* specific types of filters
- Firestone studies, 18
- Fisher studies, 127
- "Fish-finder" sonar, 70
- Fitzgerald, Stark and, studies, 211
- Fixed pattern noise, 17–18
- Fixed-point notation, 326
- Fizeau optic set, 56
- Flat-bed scanners, 29
- Flicker, 12, 137
- Flickner studies, 106
- Floating-point notation, 326
- Flook studies, 417
- Floppy disks, 103
- "Flower" image, 126–127
- Fluorescence (fluorescent microscopy)
- automatic ridge-following, 369
 - division, 270
 - imaging requirements, 73
 - noisy images, 137
 - optical sectioning, 598
 - voxel values, 626
- Focusing, 18, 38, 292–293

Foley and Van Dam studies, 190, 606
 Fourier transform infrared (FTIR) spectroscopy, 662,
see also Infrared imaging
 Fourier transforms, *see also* Frequency space
 fast Fourier transform (FFT), 278, 300
 wavelet compression, 123, 125
 Fractals
 analysis, 261–263
 compression, 126–128, 130
 measurements, feature specific, 515–519
 surfaces, 694–696
 Franke, Keller and, studies, 668
 Frank studies, 336, 582
 Frederik studies, 606
 Freeman, Abbasi-Dezfouli and, studies, 660
 Freeman and Davis studies, 521
 Freeman studies, 374
 “Freeze frame” playback, 19
 Frei and Chen algorithm, 244–246, 644–645
 Frequency space
 autocorrelation, 331–332
 basics, 277–278
 blurring, 324–327
 convolution, 314–317
 correlation, 328–331
 deconvolution, 322–324
 filtering, 297–314
 Fourier transform, 278–286
 frequencies and orientations, 286
 imaging system characteristics, 317–321
 masks, 303–309
 measurements, 286–297
 motion blur, 324–327
 orientation and spacing, 286–290
 periodic information selection, 309–314
 periodic noise isolation, 297–303
 template matching, 328–331
 texture and fractals, 295–297
 Frieden studies, 168
 Frisby studies, 5
 FTIR spectroscopy, *see* Fourier transform infrared
 (FTIR) spectroscopy
 Fu and Mui studies, 380
 Fuchs studies, 623
 Fudge, 556
 Fukunaga studies, 380, 547
 Functional parameters, 687, 691–692
 Fusion, *see* Stereology (stereoscopy)
 Fu studies, 553
 Future trends, 649–650
 Fuzzy classification, 380, *see also* Classification
 Fuzzy logic, 107, 542, 550

G

Gamut (color range), 88, 96–97
 Gaussian properties
 Difference of Gaussian (DOG), 243, 247
 Laplacian method, 228
 neighborhood averaging, 142–149
 noise reduction, 165
 periodic noise isolation, 243
 Gauss-Jordan elimination, 567

Gentile studies, 117
 Geographical information system (GIS), 100–101, 108,
 606, 610
 Geometrical distortion, 191–193
 Gersho, Wu and, studies, 130
 Gestalt theory, 4–5
 Gibbs studies, 567
 Global image measurements, *see* Measurements,
 global (scene-based)
 Global Positioning System (GPS), 108
 Global uniformity, 71
 Goh, Markiewicz and, studies, 668
 Golay, Savitsky and, studies, 148, 231
 Gonzalez, Tou and, studies, 553
 Gonzalez and Wintz studies, 277, 380
 Gonzalez and Woods studies, 73
 Gordon studies, 563
 Gouraud shading, 636, 676
 Gowayed studies, 643
 Gradient measurements, 466–469
 Graphs, human reliance on, 2
 Grasselli studies, 539
 Gratin and Meyer studies, 647
 Gravity waves, 1
 Green studies, 18
 Grey-scale, *see also* Halftoning
 basics, 23–24, 208
 brightness measurements, 476
 converting color images, 46
 Euclidean distance map (EDM), 427
 histogram equalization, 214
 images (binary), 172, 419
 masks, 391
 neighborhoods, 420
 printers, 80
 subtraction effect, 181
 Grigg studies, 58
 Grossberg studies, 551
 Grouping, 368
 Gundersen and Jensen studies, 460
 Gundersen studies, 400, 448
 Guthrie, Arlo, 79

H

Haar functions, 123
 Hader studies, 73
 Halford and Preston studies, 647
 Halftoning, *see also* Grey-scale
 defect removal, 170
 dot quality, 84
 printers, 80–83
 Ham studies, 583, 584
 Hand-held scanners, 29
 Hand studies, 527
 Hannah, Quam and, studies, 602
 Haralick, Zuniga and, studies, 602
 Haralick and Dinstein studies, 336–337
 Haralick and Shapiro studies, 380, 602
 Haralick studies, 259, 260, 346
 Hardware, printing, 93–98
 Harmonic analysis, 519–522, 649
 Hart, Duda and, studies, 491, 527

- Hartigan studies, 669
 - Hartley studies, 278
 - Hartwick, Holt and, studies, 106
 - Haykin studies, 527
 - HCP, *see* Hexagonal close packed (HCP) configuration
 - HDTV, 15, 28, 129
 - Hearn and Baker studies, 190, 606
 - Heath studies, 229
 - Hebbs studies, 552
 - Heckbert algorithm, 44
 - Heckbert studies, 117
 - Heijmans studies, 154, 249
 - Heipke studies, 660
 - Van Helden studies, 369
 - Van Helleputte studies, 654
 - Hergerl studies, 584
 - Herman, Milgram and, studies, 366
 - Herman studies, 558
 - Hexagonal close packed (HCP) configuration, 605
 - High-depth (resolution) images, 26–29, 58, 587–590
 - High-pass filters, 226–228, 298–299
 - Hildreth, Marr and, studies, 225
 - Hildreth studies, 225
 - Histograms
 - cumulative, 215
 - equalization, 184, 211–219
 - segmentation and thresholding, 335–336, 366–367
 - HLS, *see* Hue, lightness, and saturation (HLS) channels
 - Holmes studies, 324
 - Holograms, 2, 623
 - Holt and Hartwick studies, 106
 - Homomorphic filter, 302
 - Horizontal line readout, 13
 - Horn studies, 600
 - Hough spaces, 491–496
 - Hounsfield studies, 559
 - Howard, Reed and, studies, 486
 - Howard and Reed studies, 445
 - HSI, *see* Hue, saturation, and intensity (HSI) channels
 - HSV, *see* Hue, saturation, and value (HSV) channels
 - Huang, Yang and, studies, 152
 - Huang studies, 152, 603
 - Hubble telescope, 10, 310
 - Hubel studies, 5
 - Hue, lightness, and saturation (HLS) channels, 41
 - Hue, saturation, and intensity (HSI) channels
 - color median filters, 154
 - color spaces, 41–44
 - contrast expansion, 133
 - histogram equalization, 212
 - JPEG compression, 118, 120
 - RGB conversion, 480
 - segmentation and thresholding, 337–338
 - Sobel and Kirsch operators, 238
 - two-dimensional thresholds, 338–343
 - Hue, saturation, and value (HSV) channels, 41
 - Huffman coding, 113–115
 - Human vision
 - blur *vs.* noise, 325
 - boundaries, 248
 - brightness perception, 9, 23
 - cartoons and line drawings, 225
 - colors, 351
 - compression methods, 119
 - computer superiority, 76–77, 447, 527
 - CRT images, 129
 - depth perception, 600
 - describing shape, 514
 - edges, 224–225, 379
 - grey-scale response, 84, 208
 - grouping, 368
 - human superiority, 251, 527
 - information gathering, 1
 - inhibition, 226
 - JPEG standard selection, 121
 - lighting variations, 47
 - optical deception of, 90
 - shape-from-shading, 172
 - stereoscopy, 64–65
 - surfaces, 672, 693–694
 - threshold value interaction, 363
 - vs.* imaging devices, 3–4
 - Hurd, Barnsley and, studies, 126, 127
 - Hurst studies and operator, 261–263, 352
 - Hwang studies, 130
 - Hybrid parameters, 687
- I**
- Ichikawa studies, 105
 - Image enhancement
 - contrast manipulation, 208–211
 - derivatives, 229–232
 - division, 270–276
 - fractal analysis, 261–263
 - histogram equalization, 211–219
 - image math, 265–266
 - implementation, 263–265
 - Laplacian, 219–229
 - multiplication, 270–276
 - rank operations, 248–257
 - requirements, 74
 - Sobel and Kirsch operators, 232–248
 - subtracting images, 266–269
 - textures, 257–261
 - Images, processing
 - Boolean operations, 383–389
 - boundary lines, 439–442
 - closing, 410–413
 - color spaces, 42
 - cluster, 425
 - dilation, 409–410, 417–418
 - double thresholding, 405–408
 - erosion, 409–410, 417–418
 - Euclidean distance map (EDM), 425–444
 - features, 391–405
 - grey-scale images, 419
 - grouping pixels, 391–397
 - isotropy, 413–416
 - masks, 389–391
 - measurements, 417–418
 - neighborhood parameters, 420–425
 - opening, 410–413
 - skeletonization, 435–436, 442–444
 - thickening, 439–442
 - three-dimensional, 643–647

- ultimate eroded points, 434
- watershed segmentation, 429–433

Imaging, *see also* Three-dimensional imaging

- bandwidth limitations, 18–20
- cameras, 8–18, 36–39
- color images, 30–52
- electronics, 18–20
- focusing, 18
- geometries, 578–580
- grey-scale resolution, 23–24
- high-depth images, 26–29
- image types, 52–54
- imaging requirements, 71–77
- limitations, 36–39
- math, 265–266
- multiple images, 59–64
- noise, 24–26
- pixels, 20–23
- range imaging, 54–59
- representation, 373–376
- requirements, 71–77
- stereoscopy, 64–71
- types, 52–54
- video cameras, 6–7

Imaging system characteristics, 317–321

Imposed criteria, 532, *see also* Classification

Information, image reliance on, 1–5

Infrared imaging, 479, 662, 664

Ink-jet printers, 95–98

Inoue studies, 77, 476

Interferometric light microscope

- AFM comparison, 58–59
- contours, 371
- noncontacting measurements, 658
- range imaging, 55–56, 664–665

Interlace scans, 12, 16, 19, 28

Interpolation

- Bayer pattern, 15
- camera resolution, 15
- defects, correcting, 167–168
- three-dimensional data sets, 608

Ion mass spectroscopy, 668

Ishtar Terra (Venus), 632

Isoelevation lines, 671

Isotropic, uniform, random (IUR) requirement, 451, 461, 466–467

Isotropy, 413–416

J

Jahne studies, 52

Jain, Kayaalp and, studies, 601

Jain studies, 277, 621

James studies, 547

Janesick studies, 13

Jaynes studies, 170, 567

Jenkins, Briarty and, studies, 591

Jensen, Gundersen and, studies, 460

Jernot studies, 430

Johnson and Capowski studies, 591

Johnson studies, 587

Joint Photographers Expert Group (JPEG) standard, *see also* Compression

- codecs, 128

- digital camera limitations, 36–37
- storage, 118–123
- two-dimensional images, 605

Joshi and Miller studies, 324

K

Kacmarz' method, 565

Kaes studies, 647

Kak, Rosenfeld and, studies, 73, 380, 427

Kak and Slaney studies, 558

Kang studies, 93

Kanpur studies, 362

Karssemeijer studies, 535

Kass studies, 379

Kayaalp and Jain studies, 601

Kaye studies, 418

Keating studies, 628

Keller and Franke studies, 668

Keller studies, 668

Van Kempen studies, 324

Kernel operations, 140–151, 314–317

Key frames, 129–130

Keyword database searches, 105, 107–108, 111

Khadivi studies, 127

Kim studies, 392

Kingsbury, Young and, studies, 119

Kinney studies, 587

Kirsch operator, 232–248

Kittler, Divijver and, studies, 527

Kittler studies, 335

k -nearest neighbor (kNN), 545–548

Kober studies, 158

Koch snowflake, 517

Kolmogorov dimension, 517

Kriete studies, 73, 591

Kurzydowski and Ralph studies, 445

Kuwahara studies, 170

L

Legendijk and Biemond studies, 324

LaGrange multipliers, 568

Landsat images, *see also* Satellite imaging

- classification methods, 539–540
- composition map processing, 668
- compression, 112–113
- division, 270–271
- multiple, 59

Landweber iteration, 323

Language rules, natural, 109

Lan studies, 435

Lantjeoul, Beucher and, studies, 430

Lantjeoul and Beucher studies, 430

Laplacian method

- Canny filter, 247, 510
- derivatives, 229
- image enhancement, 219–229
- image math, 266
- periodic noise isolation, 298
- sharpening effect, 242
- subtraction, 243
- texture values, 347

Laser printers, color, 94

- Learning systems, 541–545
- Ledley studies, 44
- Lee studies, 335
- Lempel-Ziv(-Welch) technique, 116
- Length, 457–459, *see also* Caliper dimensions
- Lettvin studies, 551
- Leviatdi studies, 392
- Lichtenbelt studies, 73
- Light microscopy, 54–55, *see also* specific type of microscopy
- Lineberry studies, 238
- Lin studies, 620
- Liquid crystal display (LCD), 50–52, 88
- Li studies, 130
- Lobregt studies, 647
- Local equalization, 218–219
- Local variations sensitivity, 71–72
- Location measurements, 481–484
- Logan, Shepp and, studies, 573
- Lookup tables (LUTs)
 - color image reduction, 44
 - color substitution, 32
 - contrast manipulation, 208
 - high depth images, 26–27
 - histogram equalization, 217
 - homomorphic filter, 302
 - reduced color palettes, 117
 - Sobel and Kirsch operators, 235
 - surfaces, 636–637
 - textural orientation, 353
- Lossless coding, storage, 111–117
- Low-pass filters, 226–228, 298–299
- Lu, Wen and, studies, 130
- LUTs, *see* Lookup tables (LUTs)

M

- Mach bands, 225
- MacQueen studies, 669
- Macroscopic camera applications, 18
- Magnetic recordings, storage, 103–104
- Magnetic resonance images (MRIs)
 - data sets, 613
 - multiple images, 63
 - pseudo-rotation views, 644
 - reconstruction, 580
 - resectioning, 608
 - segmentation and thresholding, 336
 - three-dimensional data sets, 606
 - volume fraction, 460
 - volumetric imaging, 462
- Mahalandobis distance, 538–539, 548
- Mallat studies, 123
- Malzbender studies, 673
- Mammograms, 535–536
- Mandelbrot studies, 261, 417, 694
- DeMandolx and Davoust studies, 669
- Mann studies, 557
- Markiewicz and Goh studies, 668
- Marr and Hildreth studies, 225, 244, 247
- Marr and Poggio studies, 601
- Marr studies, 5, 225, 243
- De Marsicoi studies, 106

- Masks, 303–309, 389–391
- Massalski, Engel and, studies, 584
- Mastin studies, 165
- Maximum entropy methods, 168–170, 323, 567–568
- Maximum likelihood methods, 170–171
- McAulay studies, 123
- McCulloch-Pitts model, 551
- Measurements, feature specific
 - brightness, 475–480
 - caliper dimensions, 507–509
 - circles and ellipses, 505–506
 - counting, 496–502
 - describing shape, 512–515
 - fractal dimension, 515–519
 - harmonic analysis, 519–522
 - location, 481–484
 - neighbor relationships, 491–496
 - orientation, 484–486
 - perimeter, 509–512
 - size, 502–505
 - three-dimensional, 523–526, 647
 - topology, 522–523
- Measurements, general
 - frequency space, 286–297
 - images (binary), processing, 417–418
 - imaging requirements, 72–73, 75
 - resolution, 294
 - three-dimensional images, 647–649
- Measurements, global (scene-based)
 - anisotropy, 466–469
 - ASTM grain size, 453–456
 - basics, 445–450
 - curvature, connectivity, and disector, 464–466
 - determining number, 461–464
 - gradients, 466–469
 - length, 457–459
 - sampling strategies, 459–461
 - size distributions, 469–471
 - stereology, 445–450, 471–473
 - surfaces, 450–453, 456–457
 - three-dimensional images, 647–648
- Mecholsky and Passoja studies, 261
- Mecholsky studies, 261
- Medial axis transform (MAT), 436
- Median filters
 - color correction, 186
 - neighborhood ranking, 152–158
 - rank leveling, 180
 - rank operations, 249
 - speckle noise, 355
 - subtraction, 266
- Media storage, 101–103
- Medical imaging, *see* specific type of imaging
- Medioni and Navatia studies, 601
- Mendelsohn, Prewitt and, studies, 335
- Mercator's projection, 188
- Metallographic imaging, 595
- Metrology, 660
- Mexican hat (sombrero) operator, 247
- Michelson-Morley interferometer, 658
- Microscopic camera applications, 18
- Microstructural research, 75
- Milgram and Herman studies, 366

- Milgram studies, 194
- Miller, Joshi and, studies, 324
- Minkowski methods, 417, 517, 696–698
- Minsky and Papert studies, 551
- Mirau optic set, 56
- Mirrors (prisms), 658, 659
- Mitchell and Bonnell studies, 296
- Mode filter, 157
- Moiré patterns, 308
- Monck studies, 324
- Monitors (computer), 15–16, 50–52
- Moravec studies, 602
- Morphology (morphing), 202–205, 409, 420–425
- Morse code, 113
- Motif approach, 685
- Motion, *see* Sequential images
- Motion blur, 324–327, *see also* Blurring
- Motion flow technique, 60, 65
- Mott-Smith studies, 436
- Mott studies, 669
- Movies, digital, 128–130, 160, 205
- Moving Pictures Expert Group (MPEG) standard, 129–130, 605–606
- MP3 format, 129
- MRIs, *see* Magnetic resonance images (MRIs)
- Mui, Fu and, studies, 380
- Multiband images, 336–338
- Multiband thresholding, 341–346
- Multidimensional classification, 536–541
- Multiple images, 59–64
- Multiplication, 270–276
- Muybridge, Eadweard, 609
- Myler and Weeks studies, 73

N

- Nathan and Curlander studies, 165
- National Television Systems Committee (NTSC), 28, 33–35
- Natterer and Wubbeling studies, 558
- Natterer studies, 558
- Natural language rules, 109
- Navatia and Medoni and, studies, 601
- Neal studies, 510
- Negoita and Ralescu studies, 542
- Neighborhood techniques
 - defects, correcting, 140–167
 - fractal analysis, 261–262
 - histogram equalization, 217–219
 - isotropy, 413–416
 - measurements, 358, 491–496
 - method implementation, 263–265
 - morphology, 420–425
 - nearest neighbor measurements, 140–167
 - rank operations, 249
 - shape of edge-finding operations, 247
 - time requirements, 646
- Neural networks, 551–553
- Nevatia and Babu studies, 435
- Niblack studies, 106
- Nieminen studies, 157
- Nikolaidis and Pitas studies, 73, 644
- Nobel prize, 559

Noise

- basics, 24–26
- complementary metal oxide semiconductor (CMOS) cameras, 17–18
- defects, correcting, 137–140
- electron transfers, 13
- filtering images, 297–303
- fixed pattern noise (FPN), 17–18
- position information, 360
- scanners, 29
- spikes, 298, 303, 305
- Weiner deconvolution, 322–324
- Noncontacting measurements, 657–659
- Non-planar view defects, 188–189
- Nonuniform illumination, 171–174
- NTSC, *see* National Television Systems Committee (NTSC)
- Nuts learning systems example, 542–544
- Nyquist frequency, 279

O

- Object motion measurement, 268
- O'Callaghan studies, 360
- Oistämö and Neuvo studies, 154
- Oldham studies, 585
- Olsson studies, 247, 407
- Olympic filter, 162
- Opal phytoliths, 544–545
- Opaque properties, 465–466, 556, 598–600
- Opening (erosion/dilation), 410–413
- Optical character recognition (OCR), 527–528
- Optical density, 476
- Optical sectioning, 61, 596–598
- Orientation, 286–290, 484–486
- Otsu studies, 335, 364
- Outlining of regions, 167
- Out-of-focus, *see* Blurring
- Overshooting, 225–226

P

- Palettes, reduced color, 117–118
- PAL (European) television, 18
- Pancorbo studies, 58
- Panda and Rosenfeld studies, 349
- Pao studies, 527
- Pap smears, 535–536
- Parallax (vergence), 65
- Parametric description, features, 529–534
- Parker studies, 73, 364
- Parzen window function, 299
- Passoja, Mecholsky and, studies, 261
- Pavlidis studies, 73, 435, 553
- Pearsall studies, 545
- Peleg studies, 261
- Pentland studies, 105, 261, 547, 600, 659
- Perimeter, features, 509–512
- Periodic information selection, 309–314
- Periodic noise isolation, 297–303
- PET, *see* Positron-emission tomography (PET)
- Pfalz graph, 693

- Phong shading, 673, 676
 - Photo-CD (Kodak)
 - file storage, 101
 - JPEG comparison, 119, 120
 - multiple resolution versions, 109
 - Physical sectioning, 596
 - Phytoliths, 544–545
 - PICT standard format, 110, 116
 - Piezo device, 16
 - Piperno studies, 545
 - Pitas, Nikolaidis and, studies, 73, 644
 - Pitas studies, 73
 - Pixels, *see also* Resolution
 - basics, 20–23
 - calibration, 481
 - cameras, 8, 10, 15
 - compression, 112
 - direction information, 239–240
 - fractal analysis, 261
 - grouping, 391–397, 392
 - HDTV, 28
 - histograms, 217, 366
 - image enhancement, 74
 - imaging systems, 375, 505
 - neighborhood ranking, 152
 - printing, 79
 - rounded address values, 199
 - shuttering problems, 12
 - speckle (random noise), 320, 355
 - two-dimensional images, 207
 - voxels comparison, 62
 - Plane photography, *see* Aerial photography
 - Point operations, 134
 - Point spread function (PSF), 318, 321–322
 - Poisson distribution, 487
 - Polarized light microscopy, 272
 - Polaroid instant prints, 82
 - Polynomial curves, 508
 - Population classes, 537–541
 - “Portrait of a Lady,” 186
 - Position information, 359–366
 - Positron-emission tomography (PET), 579–580, 606, *see also* Emission tomography
 - Posterization, 216
 - PostScript language, 79, 86
 - Power spectrum (spectral density), 279, *see also* Frequency space
 - Pratt studies, 73, 277, 380
 - Predictive vector quantization, 130
 - Preferred orientation, 290–295
 - Presentation, surfaces, 669–674
 - Preston, Halford and, studies, 647
 - Prettyman studies, 573
 - Prewitt and Mendelsohn studies, 335
 - Printers, 79–83, 93–98
 - Printing
 - basics, 79–83
 - color, 88–93
 - dots, 83–88
 - hardware, 93–98
 - Processing images, *see* Images, processing
 - Profile measurements, surfaces, 683–686
 - Profilometers
 - physical contact, 654–655
 - range imaging, 54–55, 665
 - technology comparison, 28
 - Progressive scans, 16, 28
 - PSD standard format, 110
 - Pseudo-color
 - grey-scale value, 54
 - lookup tables, 211
 - multiply connected surfaces, 642–643
 - pixel value changes, 31–32
 - three-dimensional images, 615
 - PSF, *see* Point spread function (PSF)
 - Pythagorean distance, 427–428, 510
- ## Q
- QBIC, *see* Query by image content (QBIC)
 - Quality control applications, 75, 268, 660
 - Quam and Hannah studies, 602
 - Quate studies, 654
 - Query by image content (QBIC), 105–106
 - Quicktime movies, 128, 609
- ## R
- Radar imaging, *see* Synthetic aperture radar (SAR)
 - RAID, *see* Redundant hard disks (RAID)
 - Ralescu, Negoita and, studies, 542
 - Ralph, Kurzydowski and, studies, 445
 - Ramanesh studies, 15
 - Random pixel noise (speckle), 320, 355
 - Range imaging, 54–59, 664–668
 - Rank leveling, 179–181
 - Rank operations, 248–257, 258, 646
 - Ray tracing, 624–628, 673
 - Readout noise reduction, 13
 - “Real time” imaging, 25
 - Reconnaissance photos, *see* Aerial photography
 - Reconstruction, three-dimensional images, 558–567, 589
 - Red, green, blue (RGB) signals
 - color median filters, 134
 - color printing, 88–90
 - color space, 40–46
 - contrast expansion, 133
 - histogram equalization, 212
 - ink-jet printers, 96
 - reduced color palettes, 117
 - thresholding, 337–338
 - two-dimensional thresholds, 338–343
 - Reddy and Chatterji studies, 606
 - Reduced color palettes, 117–118
 - Redundant hard disks (RAID), 104
 - Reed, Howard and, studies, 445
 - Reed and Howard studies, 486
 - Reed studies, 487
 - Reeves studies, 278, 379
 - Reflection, 628–633
 - Region-growing method, 378–379
 - Reiss studies, 58
 - Remote sensing, 534, 539, 668
 - Rendering surfaces, 674–677
 - Reproducibility, 357–359

- Rescaled range analysis, 261
- Resolution, 15–16, 99, *see also* High-depth (resolution) images; Pixels
- RGB signals, *see* Red, green, blue (RGB) signals
- Rhines, Dehoff and, studies, 445
- Richardson studies and methods, 324, 516, 517, 694, 697
- Ridge-finding algorithm, 253
- Rigaut studies, 335
- Ringing reduction, 300
- Ritter and Wilson studies, 73, 435
- Ritter studies, 392
- Roberts' Cross operator
 - derivatives, 231–232
 - neighborhood implementation, 264
 - Sobel and Kirsch operators, 235, 247
- Roberts studies, 529, 600
- Robinson studies, 55
- Rock studies, 5
- Rolling ball filter, 166–167, 252, 255
- Rolston studies, 549
- Rosen and Crafoord studies, 685
- Rosenblatt studies, 551
- Rosenfeld, Panda and, studies, 349
- Rosenfeld, Weszka and, studies, 366
- Rosenfeld and Kak studies, 73, 380, 427
- Rotini noodles, 556
- Rovner, Russ and, studies, 544
- Rovner studies, 545
- Rubber-band boundary, 504, *see also* Boundaries
- Rubber-sheeting, *see* Warping
- Rui studies, 105
- Rumelhart studies, 552
- Run-length encoding (RLE)
 - lossless coding, 116
 - storage, 110, 373–374
 - touching of pixels, 392
 - two-dimensional images, 605
- Russ and Dehoff studies, 445
- Russ and Rovner studies, 544
- Russ and Russ studies
 - perimeter length change, 363
 - Pythagorean distances, 428
 - segmentation, 430
 - threshold settings, 335
 - topological properties, 649
- Russell and Batchelor studies, 654
- Russ studies
 - asbestos fiber identification, 495
 - boundary as fractal, 417
 - color median filter, 154
 - cone-beam geometry, 587
 - digital imaging processing, 73
 - erosion, 435
 - Euclidean distance map (EDM), 434
 - forgery detection, 391
 - fractal dimensions, 296, 695
 - grey scale values, 46
 - grid superimposition, 397
 - hole filling, 433
 - Hough transform, 495
 - nearest-neighbor distance method, 487
 - rescaled range analysis, 261
 - segmentation, 380

- stereological methods, 445
- surfaces, 659, 698
- threshold settings, 335
- weight values, 142

S

- Sabins studies, 77, 271, 539
- Sahoo studies, 335
- Sampling strategies, 459–461
- Sanchez and Canton studies, 73
- Sanchez and Zadeh studies, 542
- SAR, *see* Synthetic aperture radar (SAR)
- Sartor and Weeks studies, 154
- Satellite imaging, *see also* Aerial photography;
 - Astronomical imaging; Landsat images; Weather imaging
- alignment, 195
- Boolean operations, 387
- classification, 379–380
- color shading, 181, 184
- geographical information system (GIS), 100
- geometrical distortion, 191
- GOES satellite, 188
- ice floe tracking, 268
- multiple images, 60
- remote sensing classification, 534
- subtraction, 268
- texture, 259
- two-dimensional thresholding, 340
- Savitsky and Golay studies, 148–149, 231
- Scanners, 29
- Scanning electron microscopes (SEMs)
 - Boolean operations, 387
 - brightness measurements, 479
 - color imaging, 30
 - distortion, 191, 201
 - filling holes, 393
 - fractal surface dimensions, 296
 - grey-scale images, 24
 - image types, 53
 - imaging requirements, 73
 - multiband thresholding, 344
 - multiple images, 59
 - noisy images, 139–140
 - nonuniform illumination, 173
 - particle selection, 400
 - physical contact, 654
 - Polaroid instant prints, 82
 - profile measurements, 683
 - range imaging, 54
 - rank leveling, 179
 - stereoscopy, 66, 69, 600, 661
 - surfaces, 660, 681–683
 - technology comparison, 28
 - voxel values, 626
- Scanning tunneling microscopes (STMs)
 - contours, 371
 - high-resolution limit, 58
 - image types, 54
 - operational mode, 656
 - physical contact, 654
 - technology comparison, 28–29
- Schalkoff studies, 527, 553

- Schneberk studies, 573
- Schwarz and Exner studies, 418, 486
- Scientific applications, *see* Astronomical imaging
- Scott studies, 693
- Screen angles, 92
- Secondary ion mass spectrometer (SIMs)
 - color usage, 615
 - multiple images, 61, 62
 - sequential removal, 599
 - surface composition, 662
 - three-dimensional images, 592
- Section planes, arbitrary, 611–614
- Segmentation, *see also* Division; Thresholding
 - accuracy, 357–359
 - boundary lines, 367–370
 - classification problem, 379–381
 - contours, 370–373
 - edge following, 379
 - histograms, 366–367
 - image representation, 373–376
 - multiband images, 336–338
 - multiband thresholding, 341–346
 - position information, 359–366
 - region-growing method, 378–379
 - reproducibility, 357–359
 - split-and-merge method, 376–378
 - textural orientation, 352–357
 - thresholding, 333–336
 - thresholding from texture, 346–349
 - two-dimensional thresholds, 338–341
- Seismographs, 585–587, 628
- SEMs, *see* Scanning electron microscopes (SEMs)
- Sensitivity, local variations, 71–72
- Sequential images, 130, 609, *see also* Digital movies
- Sequential removal, 598–600
- Serial sections, 331, 462, 592–596
- Serra studies, 409, 420, 439
- Seul studies, 73
- Shape measurements, 512–515
- Shapiro, Haralick and, studies, 380, 602
- Shapiro studies, 487
- Sharpening filter, 266
- Shepp and Logan studies, 573
- Shih studies, 583
- Shutters, 12, 25–26
- SIMs, *see* Secondary ion mass spectrometer (SIMs)
- Single-line charge-coupled device, 27
- Single-photon emission spectroscopy (SPECT), 579–580
- Size distributions
 - measurements, global, 469–471
- Size (measurements), feature specific, 502–505
- Skeletonization
 - basics, 435–436
 - caliper dimensions, 508
 - double thresholding, 407
 - rank operations, 255
 - skiz, 438, 443, 489
 - three-dimensional, 647
- Skilling, Bryan and, studies, 170
- Skilling studies, 168
- Skiz, 438, 443, 489, *see also* Skeletonization
- Slaney, Kak and, studies, 558
- Slatter studies, 549
- Slit-island method, 697–698
- Sloan, Barnsley and, studies, 127
- Smeulders studies, 508
- Smith studies, 52, 583
- Smolka studies, 157
- Smoothing, *see also* Blurring
 - convolution, 314
 - halftone images, 308
 - neighborhood averaging, 149
 - noise reduction, 162
- Snakes, 369, 379
- Snoopy posters, 80
- Snow White*, 100
- Snyder studies, 324
- Sobel operator
 - image enhancement, 232–248
 - neighborhood implementation, 264
 - texture, 347, 352
 - three-dimensional version, 645
 - thresholding, 357
- Soille studies, 409
- Solid-state chip cameras, 8, *see also* Charge-coupled device (CCD) cameras
- Sonar images, 1–2, 70
- Sonka studies, 73
- Sound wave (ultrasound) imaging, 557–558, 628
- Space applications, *see* Astronomical imaging
- Spacing (frequency space), 286–290
- Spann, Wilson and, studies, 376
- Spatial parameters, 687
- Speckle (random pixel noise), 320, 355
- Spectral density (power spectrum), 279, *see also* Frequency space
- Spectral response, chip cameras, 9
- Split-and-merge method, 376–378, 641
- Srinivasan studies, 261
- Stanford, Leland, Jr., 609
- Stark and Fitzgerald studies, 211
- Stedman, Margaret, 656
- Steel, Bright and, studies, 250
- Stefik studies, 527
- Stereology (stereoscopy)
 - basics, 64–71
 - curvature, 464
 - IUR sampling, 461
 - line length, 396–397
 - measurements, global, 445–450, 471–473
 - scanning electron microscopes (SEMs), 66, 69, 600, 661
 - surfaces, 660–661
 - three-dimensional images, 600–604, 618–622
 - volumetric imaging, 462
- Stereo studies, 405, 465
- Stewart and Burnett studies, 585
- STMs, *see* Scanning tunneling microscopes (STMs)
- Stone studies, 88
- Storage
 - backup copies, 104
 - browsing, 108–111
 - databases, 104–108
 - digital movies, 128–130
 - files, 100–101

- film recorders, 98–100
- fractal compression, 126–128, 130
- imaging requirements, 74
- JPEG compression, 118–123
- lossless coding, 111–117
- magnetic recording, 103–104
- media, 101–103
- reduced color palettes, 117–118
- thumbnails, 108–111
- wavelet compression, 123–126
- Storer studies, 116
- Stout studies, 686
- Stroboscopic imaging, 25
- Stylus instruments, *see* Profilometers
- Subsampling, color, 33
- Subtraction
 - background, 173
 - color printing, 89
 - dynamic range usage, 181
 - image enhancement, 266–269
 - rank operations, 255
- Superimposing information, 265–266, 397, 447
- Supervised classification, 532, *see also* Classification
- Suppression operator, 253, *see also* Top hat filter
- Surfaces
 - Birmingham measurement suite, 687–693
 - change tree, 693–694
 - composition imaging, 662–664, 668–669
 - contacting measurements, 653–657
 - convex, 393
 - curvature, 464
 - data analysis, 677–683
 - emerging methods, 693–694
 - encoding elevation, 32
 - fractals, 295–296, 694–696
 - light, 658–659
 - measurements, global, 450–453, 456–457
 - microscopy, 659–662
 - Minkowski methods, 696–698
 - multiply connected, 639–643
 - noncontacting measurements, 657–659
 - presentation, 669–674
 - processing, 664–669
 - production processes, 651–653
 - profile measurements, 683–686
 - range images, 664–668
 - rendering, 674–677
 - roughness, 685, 693–696
 - slit-island method, 697–698
 - three-dimensional images, 633–643
 - topographic analysis, 693–698
 - uniformity requirements, 71
 - visualization, 669–677
- Surveillance tapes, 160
- Swing studies, 476
- Syntactical models, 553–554
- Synthetic aperture radar (SAR)
 - contours, 371, 632
 - noise-reduction, 165
 - stereoscopy, 69–70

T

- Tagged image file format (TIFF) standard, 110, 116, 664
- Talbot studies, 437
- Tape drives, 103
- Tapping mode, 655
- Taut-string boundary, 504, *see also* Boundaries
- Template matching
 - cross-correlation, 528
 - derivative operations, 244
 - feature recognition and classification, 527–529
 - frequency space, 328–331
 - neighborhoods, 420
 - time requirements, 646
- TEMs, *see* Transmission electron microscopes (TEMs) and imaging
- Textures
 - frequency space, 295–297
 - image enhancement, 257–261
 - orientation, 352–357
- Thermal Infrared Multispectral Mapper satellite, 271, *see also* Satellite imaging
- Thickening, 439–442
- 35mm film, 16, 326
- Thomas studies, 685
- Thompson, Barnard and, studies, 660
- Thompson studies, 192
- Three-dimensional imaging, *see also* Imaging
 - algebraic reconstruction, 563–567
 - beam hardening, 573–578
 - color space, 338
 - defects, 568–572
 - feature specific measurements, 523–526
 - image geometries, 578–580
 - maximum entropy, 567–568
 - multiple images, 60
 - reconstruction, 558–567, 568–572
 - requirements, 75
 - tomography, 580–590
 - visualizing, 591–650
 - volume *vs.* sections, 555–558
 - voxel resolution, 525
- Thresholding, *see also* Segmentation
 - basics, 333–336
 - double, 405–408
 - rank operations, 251, 255
 - region selection, 167
 - subtraction, 266
 - from texture, 346–349
- Thumbnails, 108–111
- TIFF standard, *see* Tagged image file format (TIFF) standard
- Time-lapse photography, 25
- Time sequence multiple images, 60
- Toggle filter, 158
- Tomography
 - high resolution, 587–590
 - reconstruction, 560, 568
 - three-dimensional images, 580–587, 591
- Top hat filter
 - feature-matching, 502
 - frequency space, 305, 308

- noise reduction, 166–167
- rank operations, 250–253
- sharpening images, 313–314
- Topology, 522–523, 649, 693–698
- Tou and Gonzalez studies, 553
- Touching (features), 429–433, 497
- Toy Story*, 100
- Transfer and readout process, 12, 13
- Transforms
 - background fitting, 177
 - compression, 304
 - discrete cosine transform (DCT), 118, 125–126
 - fast Fourier transform (FFT), 278
 - Fourier, 123, 125
 - Hough, 495
 - medial axis transform (MAT), 436
 - neighborhood averaging, 149
 - voxels, three-dimensional, 278
- Transmission electron microscopes (TEMs) and imaging
 - alignment, 491
 - color imaging, 30
 - counting particles, 502
 - image types, 52
 - nonuniform illumination, 172–173
 - peak location measurement, 287
 - periodic information selection, 309–310
 - preferred orientation, 291
 - rank leveling, 179, 180
 - serial sections, 291
 - stereoscopy, 602–603
- Tristimulus correction, 47–50
- Trussell studies, 364, 366
- Tube-type cameras, 7
- Tukey, Cooley and, studies, 278
- Tu studies, 44
- Twiss studies, 545
- Two-and-one-half dimension, 172
- Two-dimensional arrays, 27
- Two-dimensional images, *see* Imaging
- Two-dimensional thresholds, 338–341

U

- Ultimate eroded points (UEPs)
 - basics, 434
 - convex shapes, 484
 - medial axis transform (MAT), 436
 - three-dimensional skeletonization, 647
 - watershed segmentation, 430–431
- Ultrasound (sound wave) imaging, 557–558, 628
- Umabaugh studies, 73
- Undershooting, 225–226
- Underwood and Banerji studies, 261
- Underwood studies, 445
- Unfolding, 471–473, *see also*
 - Stereology (stereoscopy)
- Uniformity of surfaces, 71
- Universal Product Codes (UPCs), 527
- Unsharp masking, 27, 242–243, 251
- Unsupervised classification, 532, *see also* Classification

V

- Vannier, Wang and, studies, 587
- Variable length coding, 113
- Variations sensitivity, 71–72
- Varifocal mirror, 623
- Varnish, 186
- Verly and Delanoy studies, 257
- Vertical sectioning, 452, 467
- Video cameras
 - absolute color information, 20
 - acquiring images, 6–7, 25
 - brightness measurements, 476
 - HDTV quality, 15
 - image resolution, 15
 - noisy images, 138
 - pixels, 22
 - sensors, 27
 - voltage, 6–7, 18
- Video conferencing, 38
- Video for Windows, 128
- Video line noise, 160
- Video recorders, 36
- Vidicon, 6
- Villarrubia studies, 668
- Virtual nodes, 437–438
- Virtual reality in movies, 100
- Visibility, contrast manipulation, 210
- Visualization
 - scientific applications, 32
 - software programs, 346, 625
 - surfaces, 669–677
- Visualizing three-dimensional images
 - color, 615
 - data sets, 604–611
 - data sources, 591–592
 - display hardware, 623–624
 - future trends, 649–650
 - history, 624
 - image processing, 643–647
 - measurements, 647–649
 - optical sectioning, 596–598
 - ray tracing, 624–628
 - reflection, 628–633
 - section planes, arbitrary, 611–614
 - sequential removal, 598–600
 - serial sections, 592–596
 - stereo measurement, 600–604
 - stereo viewing, 618–622
 - surfaces, 633–643
 - volumetric display, 615–618
- Voltage, 6–7, 23
- Volume elements, *see* Voxels (volume elements)
- Volume fraction, 448
- Volumetric display, 615–618
- Volume vs. sections, 555–558
- Voronoi tessellation, 203, 489
- Voxels (volume elements)
 - algebraic reconstruction technique (ART), 566–567, 583
 - boundaries, 379
 - multiple images, 62–63
 - pixels comparison, 62

three-dimensional images, 524, 525, 604–605
three-dimensional transforms, 278

W

Van der Waals forces, 655
Wall studies, 362
Walters studies, 549
Wang and Vannier studies, 587
Wang studies, 581, 660
Warping
 alignment, 194
 imaging requirements, 73
 interpolation, 200
 morphing, 202
 square pixels, 505
 subtraction, 268
Watershed segmentation, 429–433, 500
Wavelet compression, 123–126
Weather imaging, *see also* Satellite imaging
 data sets, 609
 maps, 2
 multiple images, 60
 nonuniform illumination, 172
Weeks, Myler and, studies, 73
Weeks, Sartor and, studies, 154
Weeks studies, 73
Weibel studies, 445, 471
Welch and Hanning window functions, 299
Welstead studies, 123
Wen and Lu studies, 130
West studies, 606
Weszka and Rosenfeld studies, 366
Weszka studies, 259, 335, 362
Whitehouse studies, 685
Wickramasinghe studies, 56, 654
Wiener methods, 322–324, 620
Wilson, Ritter and, studies, 73, 435

Wilson and Spann studies, 376
Winstanley studies, 549
Wintz, Gonzalez and, studies, 277, 380
Wolf studies, 335
Wong studies, 660
Woodham studies, 600
Woods, Gonzalez and, studies, 73
Wrobel studies, 660
Wu and Gersho studies, 130
Wubbeling, Natterer and, studies, 558

X

X-ray energy spectroscopy, 668
X-rays, 573–574
X-ray tomography, 63, 558

Y

Yager studies, 364
Yakimovsky studies, 377
Yang and Huang studies, 152
Yang studies, 488
YIQ/YUV color, 40, 117, 212
York studies, 557
Yoshitaka studies, 105
Young and Kingsbury studies, 119

Z

Zadeh, Sanchez and, studies, 542
Zadeh studies, 542
Zhdan, Castle and, studies, 654
Zhou and Dorrer studies, 660
Zimmerman studies, 542
Zip disks, 103
Zucker studies, 379
Zuniga and Haralick studies, 602

The Image Processing Tool Kit CD-ROM

This indispensable set of Photoshop™ compatible plug-ins delivers an impressive array of tools implementing methods in the Handbook to help you process and measure your images. It gives the familiar environment of Photoshop the analytical power of dedicated programs costing thousands more.

- Methods for correcting imaging defects
- Selective image enhancement
- Segmentation
- Morphological processing of binary images
- Global and feature-specific measurements
- Fourier image processing
- Comprehensive manual
- On-line technical support
- Macintosh and Windows compatibility

For more information go to <http://www.ReindeerGraphics.com>

FEATURES

- Provides balanced and complete access to frequency space methods with easy-to-use tools
- Presents comprehensive selection of different tools for image processing operations
- Requires only a modest technical background in mathematics or computers

CONTENTS

- > **175 Plug-ins**
Tools for everything from image adjustment and processing to measurement and rendering
- > **300 Test Images**, including many from the Handbook (plus use your own images)
- > **500 Page Electronic Tutorial** that complements the Handbook

\$50.00 off SPECIAL OFFER†

†MSRP US\$249.95. The Image Processing Tool Kit is US\$199.95 with this special offer.

Name _____
Please print clearly

Company/Institution _____

Address _____

City _____ State _____ Zip _____

Country _____

Signature _____ Date _____

Telephone No. () _____

FedEx Number (For Faster Shipping) _____

ON ALL ORDERS **SIGNATURE AND TELEPHONE NO. REQUIRED**

☐ Please send me _____ copy(ies) of
**THE IMAGE PROCESSING
TOOL KIT CD-ROM**
Catalog no. IPTK-CD at \$199.95
(MSRP US\$249.95) each. Shipping via
First Class Mail is provided. Add 6.5%
sales tax in North Carolina.

☐ Visa ☐ Mastercard ☐ Check

Credit Card Number

□ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □

Exp. Date (M/Y)

□ □ □ □

Verification code:

On the backside of your credit card, in the signature area there is a 3-digit number. Please provide it here: □ □ □

V-Code

ALL ORDERS:

Prices quoted are in U.S. dollars and are subject to change. Shipping via First Class Mail (Air Mail) is provided. If you require faster shipping, please provide your FedEx number with the order.

Order now and save \$50.00 off the regular price.

*Special offer for the
Image Processing Tool Kit CD-ROM*

For more information on the software,

*visit us at our
web site!*

<http://www.ReindeerGraphics.com>

90-DAY MATERIALS WARRANTY: Reindeer Games, Inc. makes no Warranty as to the fitness of the software for your particular task. Reindeer Games, Inc. DOES warrant the materials from defects for a period of 90-Days after purchase and will replace any defective product in that period of time. No returns except for defective materials requiring replacement will be accepted.

Mail or FAX your order to:

IPTK-CD SPECIAL ORDER, 4th ED.

REINDEER GAMES, INC.

P. O. BOX 37188

Raleigh, NC 27627

FAX: 828/252-7516, *for faster service*

**Electrical Engineering
Computer Graphics**

"Every person who uses computer image processing...who is planning to use image processing...who is involved in the purchase of computer systems or software that involve image processing...should read this book. And if you are not in one of the above groups, you might want to read it anyway."

— Microscopy Research and Technique
(on the third edition)




First published in 1992, *The Image Processing Handbook* not only set the standard for professional references in this field, but also provided the first text truly accessible to undergraduate students and non-specialists. Each subsequent edition has reflected the continuing rapid advances in image processing, and the fourth edition is no exception. In fact, its broader scope, increased depth, and logical presentation make the fourth edition, more than ever, a powerful teaching tool, essential reference, and handy companion for both new and experienced image processors and analysts.

Highlights of the Fourth Edition

- **Digital cameras and printers**—Extensive discussion of their characteristics and influence on image processing
- **Image restoration and deblurring**—A new section providing user-oriented information on how to do it well
- **Stereology**—A new chapter that makes clear the links between stereology and automatic image analysis
- **Feature classification and identification**—A new chapter that demystifies the subject and compares expert systems, fuzzy logic, neural nets, learning methods, and syntactical description
- **Almost 2000 illustrations, most in full color!**

Virtually every chapter has been updated, many quite extensively, and a multitude of new images further enhance the outstanding teach-by-example presentation. The fourth edition will undoubtedly continue this handbook's reign as the premier resource for the field of image processing.

The
**IMAGE
PROCESSING**
Handbook
Fourth Edition



CRC PRESS
www.crcpress.com

1142

ISBN 0-8493-1142-X



9 780849 311420