

Processing and Quantifying Images

CHAPTER PREVIEW

In this chapter we will equate processing with the use of the computer to analyze our data. We will simply use image processing to extract more information from the data than we can obtain by eye. The data will generally be an HRTEM image but could be other images or DPs. We'll quantify spectra after we describe them in Part 4 of this text. In the past, the optical bench was also used for this purpose, but the number of optical benches is negligible compared to the number of computers now found in every TEM lab. Optical benches did allow us to form DPs which we could then modify to produce a processed image. This analog approach has now largely been replaced by its digital counterpart. The computer can be much cheaper than the optical bench and is far more flexible. The number of software packages which are designed for, or can easily be adapted to, TEM is also growing.

We can use image processing to produce a clearer view of the image, for example, by subtracting unwanted background detail, correcting for noise or drift, or removing artifacts. The big warning, though, is that, when removing one artifact, you must be very careful not to introduce others.

Although it's nice to see information more clearly, the unique feature of the computer approach is that we can *quantify* the data in any image and then normalize these data. Now we can directly compare the quantified experimental image with computer-simulated images. Although throughout this chapter we will be concerned with HRTEM images, most of what we say can be transferred directly to the analysis of diffraction-contrast images.

The other general point is that the ideas we'll discuss are also applicable to images derived from different sources. Once the data are in the computer, i.e., in digital form, the source becomes unimportant as far as processing possibilities are concerned. Examples of 'images' which might be obtained from the TEM include X-ray or EELS maps, STEM images, TEM images, and CBED or BSE patterns.

Most of the discussion will concern the use of computers. All you need to know is how best to get the data into the computer, how to process it, what to do with the data, how to display the result, and how to report what you've done. Most importantly, when we have the image in a numerical form we can start to consider statistics, which means much more than adding 'error' bars.

31.1 WHAT IS IMAGE PROCESSING?

Image processing is essentially manipulating images. The topic arises in many fields so we need to understand the words/jargon; we'll discuss the language of image processing as it is applied to TEM.

Image processing is not only becoming more common, but is also finding new applications in many fields. Faster, more powerful computers and increased memory storage are making tasks possible which could not previously have been considered. As a result of this increasing user base, there are now

many software packages available which can be used in microscopy; we listed some in Section 1.5. These range from programs used widely in desktop publishing to those which have been custom designed for EM. The goals of image processing include that of quantitative microscopy. You must choose among the different packages, commercial and freeware, and match them to the computer available in your lab. One point to remember is that some very simple optical methods which don't rely on a computer can often be very helpful. The other point is that the eye is hard to beat.

THE BASIC IDEA

Change images into numbers and manipulate the numbers.

There are many specialized books on this topic for the beginner or the expert; a selection is given in the references. The purpose of this chapter is to give a generalized overview. One problem in discussing this topic is that it is a very rapidly changing field. We will try to avoid specifics concerning particular programs but will mention these programs at the end of this chapter.

31.2 PROCESSING AND QUANTIFYING IMAGES

We process images primarily for two reasons

- We may want to improve the appearance of an image, make it look sharper, more even in contrast, higher contrast, etc. This processing can be unwise.
- We may want to quantify the information contained in the image. This processing may be essential: many physicists can read equations but not images.

Processing for improving the appearance of TEM images has been practiced for many years using such photographic techniques as ‘dodging,’ using ‘filters,’ selecting different emulsions or varying the developer, etc. It is only recently that relatively powerful personal computers have become widely available, but the term ‘image processing’ almost automatically implies the use of computers. Computer image processing will be the emphasis of our discussion. We have three requirements.

- We must be able to create a digital form of the image in the computer.
- We need appropriate software for processing the image.
- We need a computer which can perform the processing in an acceptable period of time with the required resolution.

THE DIFFICULTY

When image processing, the difficulty may be describing what you have done.

Many comments here are similar to those we made in discussing the microscope itself. For example, you may have to work with the built-in system or the system that’s already available in the lab. The difference is that some of the free programs are extremely powerful, so that all you need is the desktop computer. Many programs designed for desktop publishing are relatively inexpensive. Thus, you can almost always find a way to extend your processing capabilities.

The motivation is that we need to obtain more information from an image than we can get by just looking at it. This principle applies to more than HRTEM; we are discussing it here because HRTEM is where at present it is most needed/used in TEM. However, any TEM, X-ray map, or energy-filtered image or DP may benefit from processing and quantification. We need to quantify the TEM parameters, in particular C_s . One unique aspect of image processing in the TEM is that we have a choice between on-line and off-line processing. In fact, we often use on-line processing (frame averaging and background subtraction on the video image) to see the image even though the image we record may not be processed. However, it may indeed already be processed—so you should know and report this.

31.3 A CAUTIONARY NOTE

For most of our discussion we will consider only processing techniques using computers. To a large extent we can simulate a TEM using the computer. As we saw in Chapter 30, we can model a crystal, insert apertures, define the electron beam, including its broadening in the specimen, and then calculate the image. What we do in image processing is start with the image, add apertures and special filters, and then create a new image, the processed image. This image is a *real* image. What we must be careful about is explaining just what processing procedures we have used, since these may affect the interpretation of the data. Reporting the processing procedure is particularly critical when the raw data (the ‘original’ image) is not being reported at the same time.

ALWAYS

Report how you have treated your image, so that the reader can compare your data with related data that may have been processed differently or not at all. Be very careful. One honest (or dishonest) mistake may change your career!

31.4 IMAGE INPUT

There are several methods you can use to put the TEM image on the computer. The choice depends in part on how much detail you want in your digitized image, but also depends on how much work you’re prepared to do. In this discussion, we’ll only consider images which you have looked at on a video monitor, a computer display, or the fluorescent screen. Your basic choices are

- Transfer the image directly from the TEM to the computer.
- Record the image on film, then digitize it using a microdensitometer.

- Record the image on videotape.
- Record the image on film, then print it and use a flat-bed (or other) scanner.

There are many methods for creating a digital form of an image in your computer. The simplest is to use a slow-scan CCD camera, which we discussed in Chapter 7. The drawback of CCD cameras is that high-quality CCD chips are very expensive for $2\text{ k} \times 2\text{ k}$ arrays and astronomically expensive for $4\text{ k} \times 4\text{ k}$ arrays (but they are getting cheaper). Such cameras may become routine additions for all TEMs in the future. You will probably also want to use video. With film, you can record a larger area than you can using a CCD unless you have the $4\text{ k} \times 4\text{ k}$ version. You should use a digital video-recorder for in situ studies when using a heating or straining holder.

We can transfer the image from a videotape or a video camera to the computer using a frame grabber. Frame-grabber boards are readily available for most computers. You can use a high-resolution scanner for photographs or negatives. At this time, scanners cost about the same as a digital camera with comparable resolution. The purist's approach is to use a microdensitometer to measure the intensity of the film point by point and read this directly into the computer. The advantage of the microdensitometer is that it is very precise and can achieve the highest resolution for a very large area. The main problem is that it is slow, being a serial-collection technique. If you use it to its best advantage, your image will require a large amount of computer memory, which in itself is not a problem, but manipulating such images will still be slow.

31.5 PROCESSING TECHNIQUES

31.5.A Fourier Filtering and Reconstruction

The principle involved in filtering is that a mask is used to remove some information from an image in order to enhance or emphasize other information. As an extra complication we can process the image, e.g., Fourier transform a HRTEM image, then apply a mask, and then reverse the processing.

We can vary the size of the apertures and the sharpness of their edges, something you can't do in a TEM with normal fixed-diameter objective apertures. A single variable SAD aperture was used on some early TEMs; it was triangular or square in shape and used three or two movable blades. You can best understand the procedure by an example. A square mask was used to select the region in Figure 31.1A from a much larger region of the HRTEM image, and its Fourier transform (i.e., the DP effectively from a few nanometers) is shown in Figure 31.1B.

THE SUPPORT FILM IS IMPORTANT

You will often see similar diffractograms obtained using a film of amorphous carbon. While such films are easier to make, they give little diffracted intensity for the range of u values between 6 and 8.5 nm^{-1} , which is important in HRTEM.

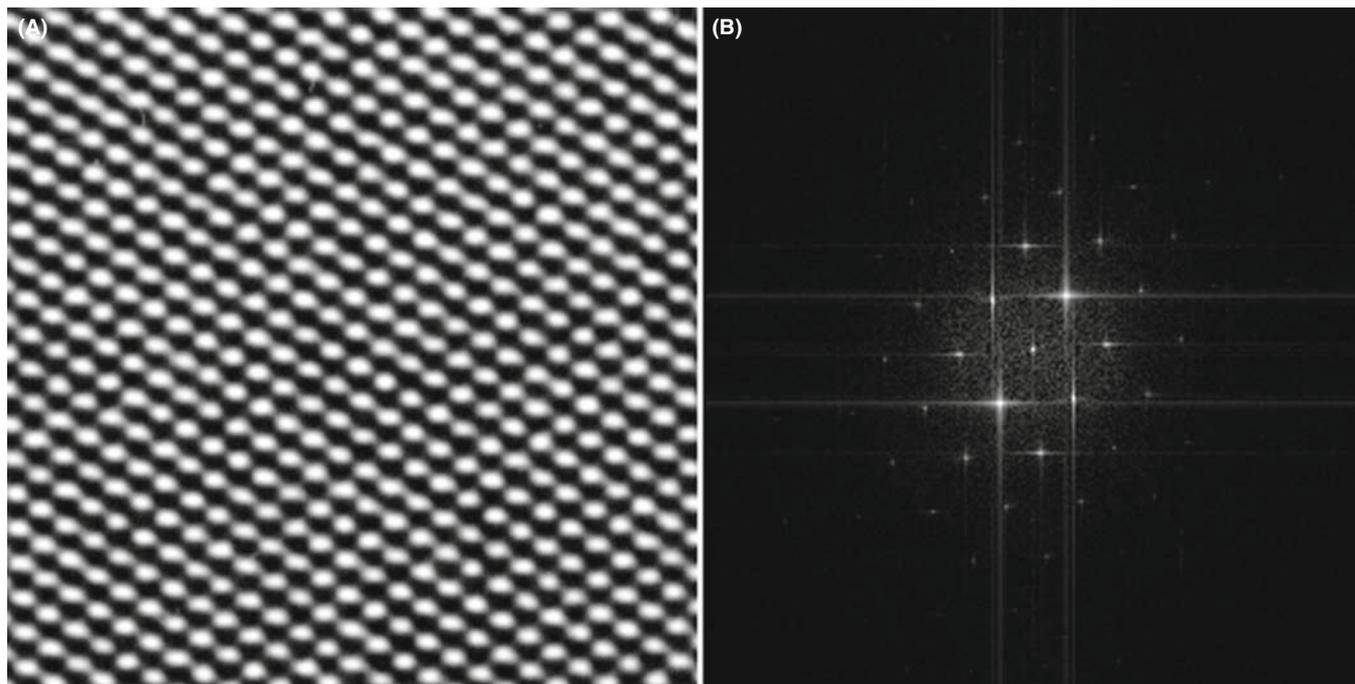


FIGURE 31.1. A square mask has been used to select the area shown in (A) from a much larger print of the image. The Fourier transform of this region is shown in (B) where you can see not only the spots in the 110 DP but also long streaks (artifacts of the processing) that run normal to the edges of the mask.

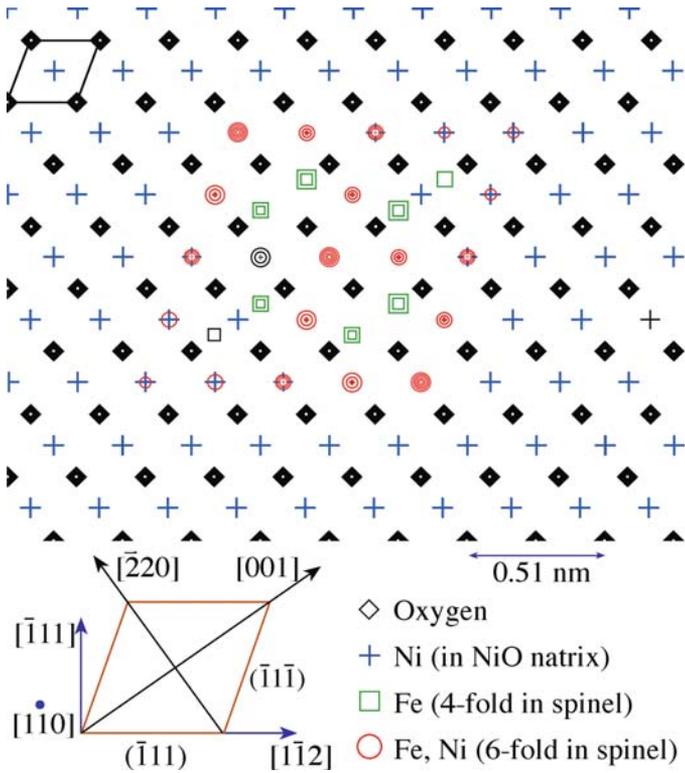


FIGURE 31.2. A model of an octahedron of spinel fully enclosed within a matrix of NiO. The rest of the specimen could then be modeled by adding extra layers of NiO above or below the defect layer.

What this technique does is to allow you to do microscopy in the computer. Your image becomes the specimen. You form the DP, then you can use apertures to select one or more beams to form the image; these apertures are the computer version of the objective aperture. Small apertures limit resolution just as in the 'real' TEM because the information about anything other than the perfect lattice is carried between the reciprocal lattice points. Figure 31.2 gives an illustration of how a model can be constructed of a particle in a matrix, which can be useful when simulating HRTEM and conventional BF/DF images in the computer, as shown in Figure 31.3. This model was made using the Digital Micrograph package (see Section 1.6).

31.5.B Analyzing Diffractograms

In Chapter 28 we showed that the transfer function could be plotted out as shown schematically in, e.g., Figure 28.4. Another way of thinking about this plot is to imagine what would happen if we have a specimen which generated equally every possible value of \mathbf{u} , i.e., every possible spatial frequency.

An amorphous film of Ge can provide just such a plot, but it is difficult to record the result because the scattered intensity is low.

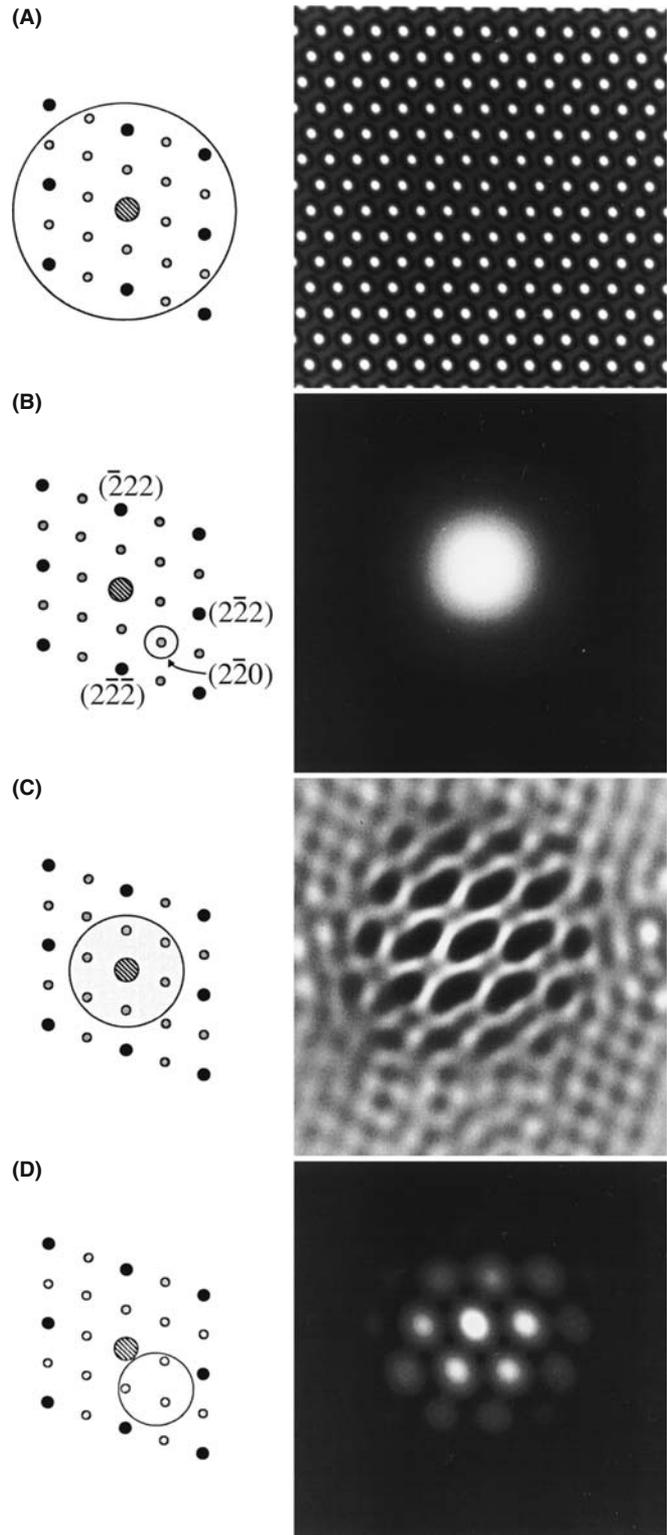


FIGURE 31.3. The DP from a model specimen containing a particle such as that shown in Figure 31.2 is shown schematically in (A) together with the resulting lattice image. The other three pairs of diagrams illustrate how we can use the computer to produce different masks and thus generate different images, such as the DF image in (B). The image in (C) corresponds to the image shown in Figure 28.26, while the DF lattice image shown in (D) is analogous to that discussed in Figures 28.14 and 28.15.

We therefore record the image at high resolution, preferably directly, using a slow-scan CCD camera, although digitizing the negative is fine. Then by comparing the experimental plot of I versus \mathbf{u} with those calculated for different values of Δf and C_s we can determine the astigmatism, the defocus, Δf , and the value of C_s (as we'll see below). It helps if you have a few particles of Au on the Ge film since the Au spots then give an internal calibration. Such a set of images and their corresponding diffractograms is shown in Figure 31.4. Notice that as the defocus of the objective lens increases the number of rings increase but they become narrower. The contrast transfer gradually extends to larger values of \mathbf{u} .

Determining astigmatism. you can use such diffractograms to correct the astigmatism, since a perfectly stigmatized image will give a DP with circular symmetry. As you can see in Figure 31.5, even a small amount of astigmatism can be detected by eye. The computer can readily measure and provide feedback to the TEM to correct the lens defect, as we will see shortly. This set of diffractograms also shows that the computer can distinguish astigmatism and drift in the image while the eye can easily mistake one for the other. Drift produces a circular pattern but the higher spatial frequencies are lost in the direction of drift.

Determining Δf and C_s . You can determine Δf for any image by measuring the radii of the bright and dark rings in the diffractogram, since bright rings correspond to $\sin \chi(\mathbf{u}) = 1$ and dark rings correspond to $\sin \chi(\mathbf{u}) = 0$.

$$\sin \chi(\mathbf{u}) = 1 \quad \text{when } \chi(\mathbf{u}) = \frac{n\pi}{2} \text{ and } n \text{ is odd} \quad (31.1)$$

$$\sin \chi(\mathbf{u}) = 0 \quad \text{when } \chi(\mathbf{u}) = \frac{n\pi}{2} \text{ and } n \text{ is even} \quad (31.2)$$

Since C_s will also influence the location of the rings, you need at least two rings. Krivanek has given a simple procedure for finding both C_s and Δf . If we start with our definition of χ

$$\chi(\mathbf{u}) = \pi\Delta f\lambda u^2 + \frac{1}{2}\pi C_s\lambda^3 u^4 \quad (31.3)$$

then, inserting the values given in equations 31.1 and 31.2 leads to

$$\frac{n}{u^2} = C_s\lambda^3 u^2 + 2\Delta f\lambda \quad (31.4)$$

All we now have to do is plot nu^{-2} versus u^2 to obtain a straight line with slope $C_s\lambda^3$ and with an intercept on the nu^{-2} axis of $2\Delta f\lambda$. Assign $n=1$ to the intensity maximum of the central bright ring, $n=2$ to the first dark ring, etc. The analysis can be trickier if you have used an underfocus condition or if you are very close to Scherzer

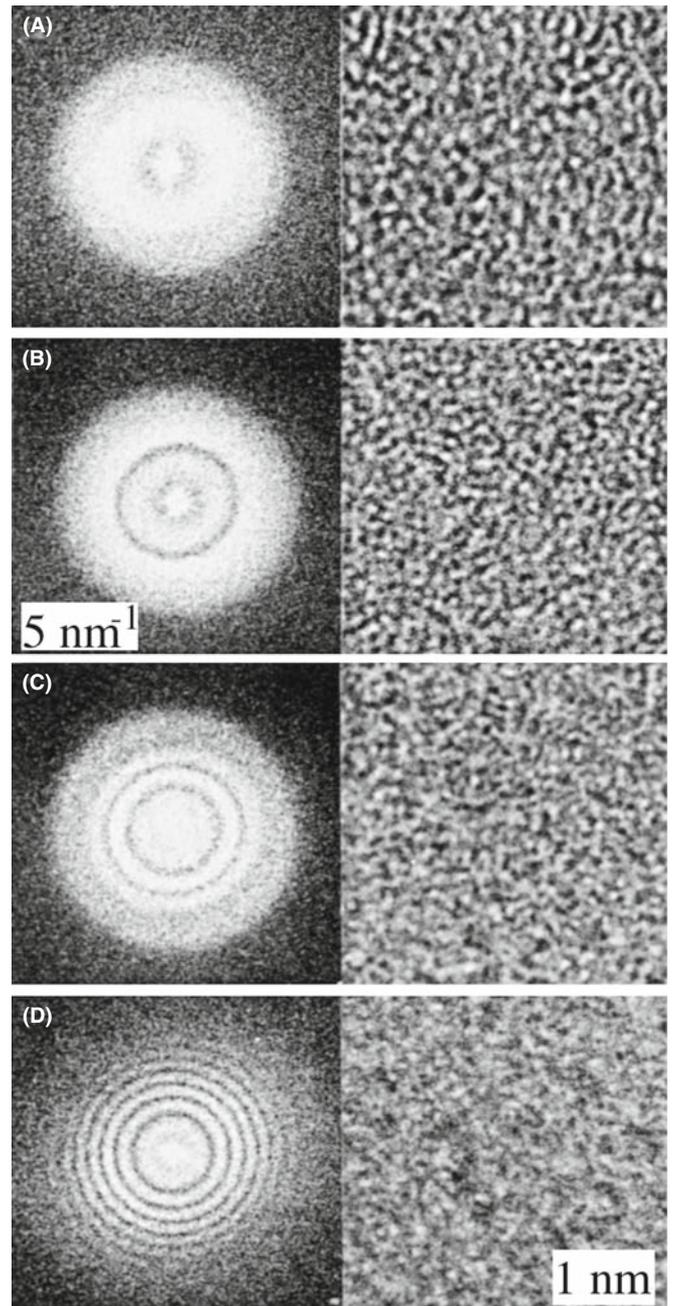


FIGURE 31.4. Four images of an amorphous Ge film and their corresponding diffractograms. Δf has the following values: (A) 1 sch; (B) 1.87 sch; (C) 2.35 sch; (D) 3.87 sch. Remember that 1 sch = $-(C_s\lambda)^{1/2}$.

defocus, but you will know when you have not found a straight line! Hopefully, you'll find that your value of C_s is close to that given by the manufacturer! A rather neat result is that if you plot nu^{-2} versus u^2 for different diffractograms (i.e., different values of Δf), then the points corresponding to each particular value of n will lie on a hyperbola, as shown in Figure 31.6A. You can use these hyperbola to determine C_s for any microscope and Δf for any diffractogram.

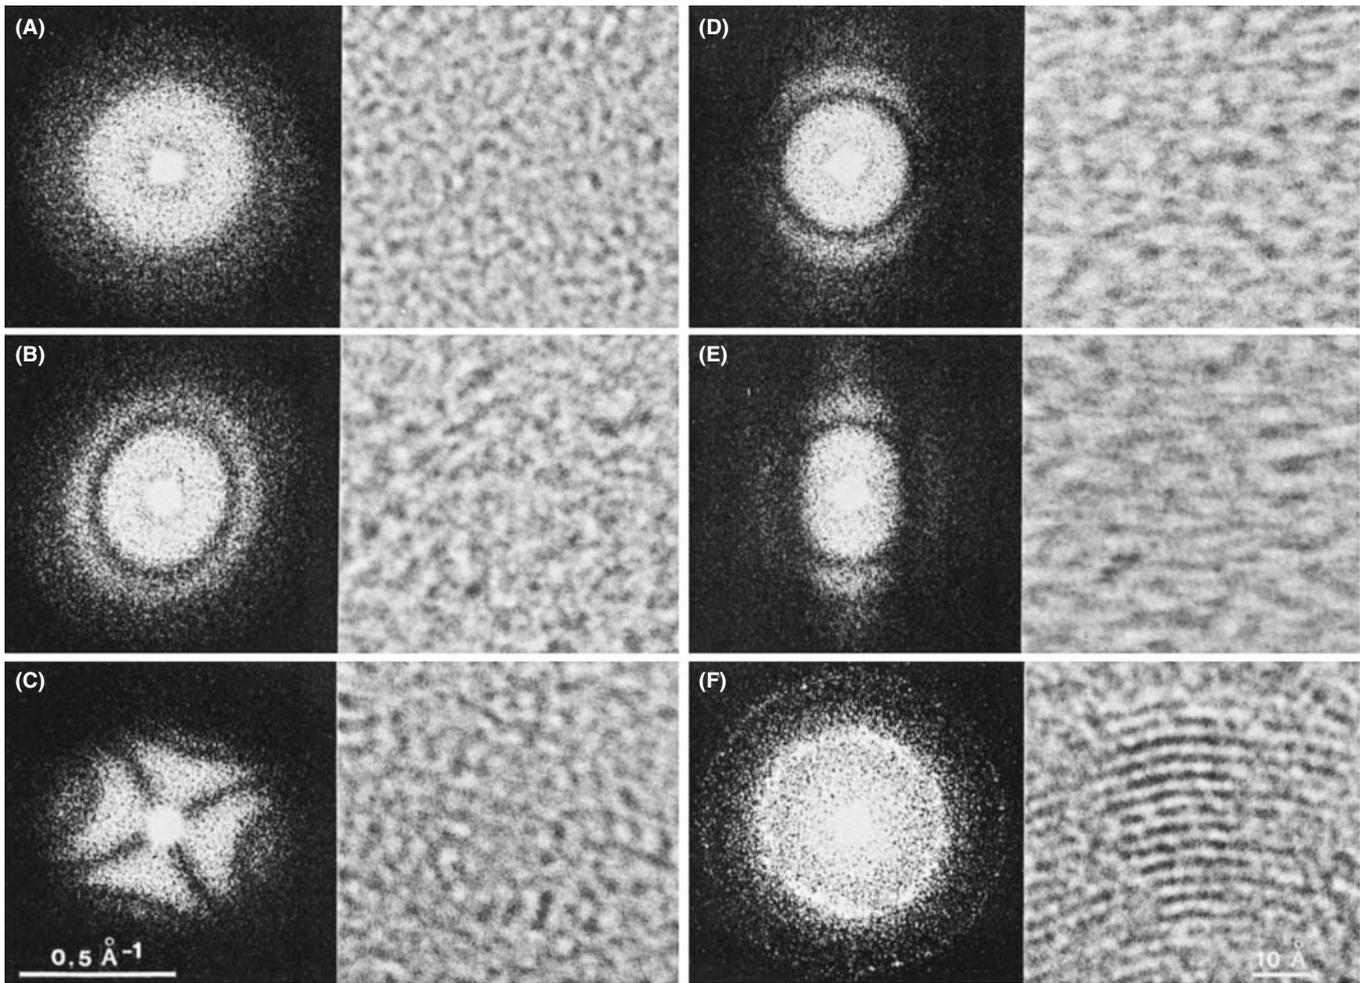


FIGURE 31.5. Six images of an amorphous carbon film and their corresponding diffractograms illustrating different misalignments of a 300-kV HRTEM. (A) Well aligned and no drift; (B) some astigmatism ($C_a = 14$ nm); (C) more astigmatism ($C_a = 80$ nm); (D) no astigmatism but drifted 0.3 nm; (E) no astigmatism but drifted 0.5 nm; (F) well aligned and no drift showing graphite calibration fringes of 0.344 nm spacing. (B,D,E) $\Delta f = 2.24$ sch; (F) $\Delta f = 0$.

THE ZEMLIN TABLEAU

This is a display of diffractograms recorded by systematically changing a parameter, such as tilt or defocus, as in Figure 31.5.

Diffractograms and beam tilt. Beam tilt is very difficult to correct by eye; even worse, it can cause the diffractogram to look astigmatic, so you correct the astigmatism instead. In the image, as we saw earlier, beam tilt can improve the appearance but confuse the interpretation! The set of diffractograms shown in Figure 31.6B shows you how to overcome the problem. You have to compare diffractograms taken at different beam tilts to determine the zero-tilt condition. A pair of diffractograms taken at $\pm\theta^\circ$ tilt will only look the same (though rotated) if the beam had zero tilt at $\theta=0^\circ$. In the example shown, the diffractograms above and below the horizontal line are similar, so θ_y was very close to zero for the central condition.

However, the pairs of diffractograms on opposite sides of the vertical axis differ slightly, so the alignment of θ_x was not perfect.

31.5.C Averaging Images and Other Techniques

If you have recorded a series of images using a video camera, for example, you can average them over several frames as your eye does automatically. The result of such a process is illustrated in Figure 31.7. Different methods can be used to average the images. The easiest approach appears to be as good as any and simply involves taking the unweighted average of your best images, i.e., in the video example, just average over a series of frames. If you know that the object you're studying has a certain symmetry, you can use that information to improve the image further. The article by Trus et al. will give you a start on this process. If you want to remove the blur due to motion of the image, then you will really need to delve much more into this subject.

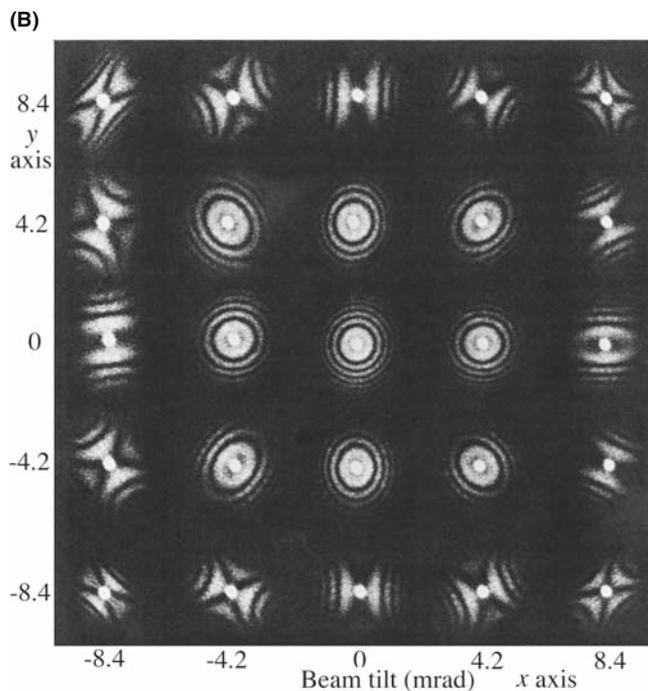
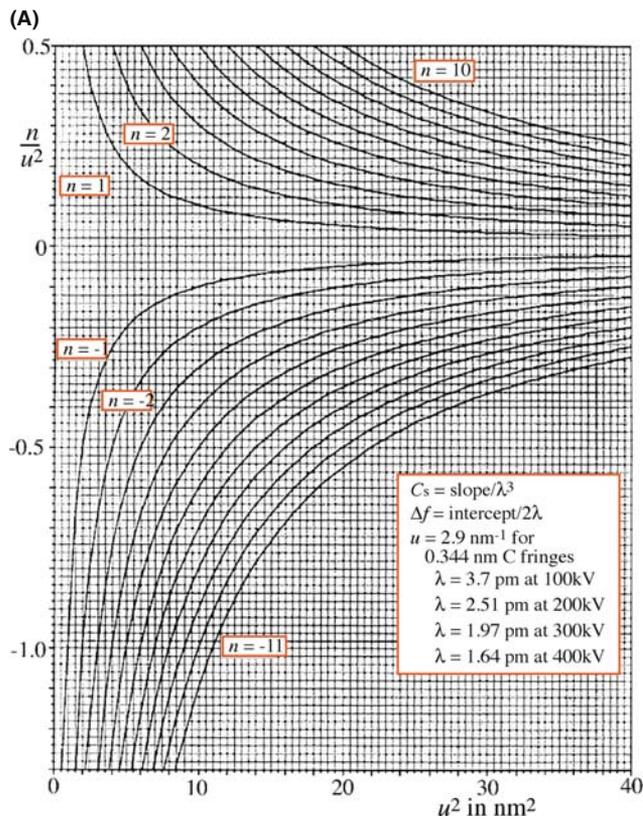


FIGURE 31.6. (A) Plot of nu^2 versus u^2 . The rings in any diffractogram correspond to a series of n values which allow you to draw straight lines on this figure and thus determine the slope and the ordinate intercept, giving C_s and Δf , respectively. (B) Set of diffractograms showing the effect of the incident beam tilt.

If you use a TV-rate video, you'll almost certainly use background subtraction routinely. For example, you can record an image of the honeycomb

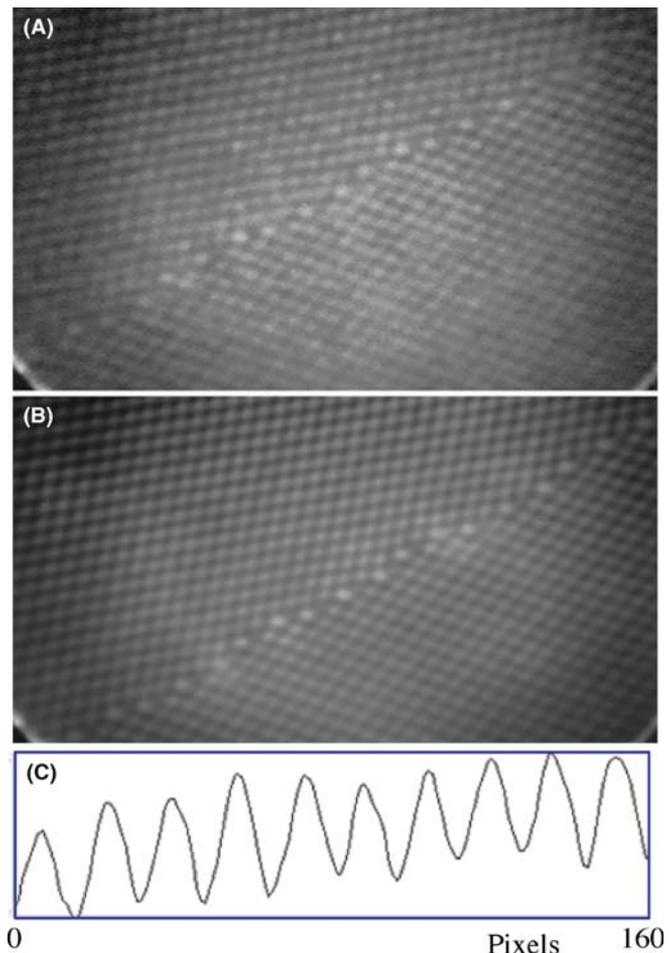


FIGURE 31.7. An example of the benefit of frame averaging to improve information from a video recording: (A) one frame; (B) 16 frames; (C) intensity profile along a (111) plane in (B).

pattern of the YAG detector, store it, and then automatically subtract it from all subsequent images in real time.

You may find it useful to add artificial color (pseudo-color) to your TEM images, as illustrated in Figure 31.8. Although it is often assumed that this is done just to make the images even more appealing to the non-microscopist (or non-scientist), there is actually a valid reason for the practice. Our eyes are much more sensitive to small variations in color than they are to small variations in gray level. You might therefore find color useful if you have a wide range of gray levels and want the viewer to be able to 'see' some subtle variations. Similarly, you can use color to emphasize a particular gray level in an image. However, you have to be very careful in your choice of look-up table (LUT), the table which relates each gray level to a particular color. To get a feel for the dangers, play with Photoshop™ and your favorite TEM image.

As we've already noted, in TEM, all of our apertures have relatively sharp edges but in the computer, you have the possibility of using multiple apertures, apertures with different shapes and apertures with diffuse

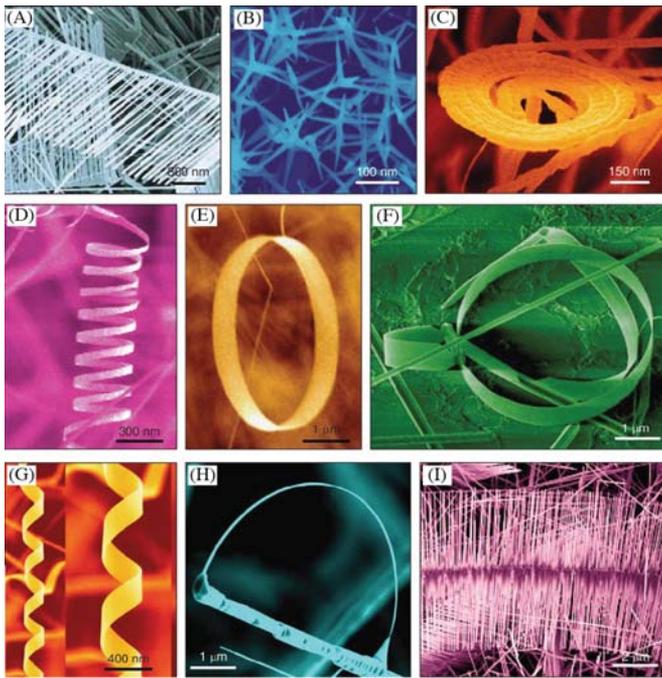


FIGURE 31.8. Illustrating the effectiveness of using false color.

edges. Apertures with diffuse edges will help eliminate the streaking which will otherwise be present (go back and look at Figure 31.1). You can also use the computer to do ‘unsharp masking,’ which is not the same as simply using a diffuse mask. The technique comes from the photographic process whereby we first print an image out of focus onto film, thus making a complementary image except where there is fine detail present in the original image; in digital processing this is called Laplacian filtering. Many more examples are given in Russ’s two books.

31.5.D Kernels

A kernel is simply an array of numbers which we can use to perform operations on a digital image. If we have the 3×3 kernel, K (we can have 5×5 , 7×7 , etc., but the computation time becomes too long, especially for real-time situations)

$$K = \begin{matrix} -1 & -1 & -1 \\ -1 & +8 & -1 \\ -1 & -1 & -1 \end{matrix} \quad K_o = \begin{matrix} A & B & C \\ D & E & F \\ G & H & I \end{matrix}$$

we can apply it to every 3×3 group of pixels in our image, e.g., K_o and put the result in a new digital image. If we call our new 3×3 image, K_i then

$$K_i = \begin{matrix} A' & B' & C' \\ D' & E' & F' \\ G' & H' & I' \end{matrix}$$

The new image will have, for example, $E' = 8E - A - D - G - B - H - C - F - I$. This kernel then gives us a digital Laplacian (an approximation to the second linear derivative, ∇^2). What this kernel is doing is subtracting the brightness value of each neighboring pixel from the center pixel. If the area is a uniform gray it will become white so changes in contrast will be exaggerated. We can design a wide range of kernel operators. For example, the edge enhancer kernel has the effect of digitally differentiating the image. (We’ll see a related digital-processing procedure applied to spectra in Chapters 35 and 39.) The Sobel and Kirsch operators are examples of such edge detectors; each can be thought of as the sum of several kernel operators. We can also use binary morphological operators which make binary features become larger or smaller. All of these operations can be carried out in any standard image-processing package. In general, you should be very careful when using such techniques in TEM; their value is in displaying data which might otherwise be missed, rather than helping you quantify an image.

REPORT IT

We’ve said it before. Check HREM papers in the literature. Have the images been processed? Are you told how? Can you reproduce the ‘experiment’? The last test is the essential component of science.

31.6 APPLICATIONS

This section will give you a taste of how image processing is being used now. It is just part of a rapidly growing list, so we are not going to be detailed or inclusive. We can separate the applications into two groups

- Noise reduction or improving the signal/noise ratio.
- Quantifying images.

Of course, the first topic is included in the second.

31.6.A Beam-Sensitive Materials

Low-dose microscopy necessarily implies that the signal-to-noise ratio will not be large; if it is large, the dose could have been smaller. This problem has been extensively addressed in biological EM and led to Klug’s Nobel prize for ‘Development of crystallographic electron microscopy and the structural elucidation of biologically important nucleic acid–protein complexes’ in 1982 (see the 1971 paper by Erickson and Klug). In materials science, we have tended to accept ‘beam damage’ as a fact of life, but this attitude will not be acceptable for future quantitative HRTEM. Most modern microscopes will allow you to perform all your

alignments on one area and then translate the beam to a predetermined distance in a predetermined direction before recording the image of a pristine area. Clearly, the CCD camera will not only let you see your image without waiting to develop the plates, but you can take a series of images for noise reduction purposes and/or assess whether the imaging conditions were what you had intended. (Sometimes, by taking a chronological sequence of images you can get a sense of the 'undamaged' image by extrapolating information back to zero time). The image shown in Figure 31.9 illustrates the possibilities. If you read the review by van Heel et al., you will get some idea of how far you can already go in this field.

BEAM DAMAGE

The electron beam will invariably change your specimen. Historically, we have worried about beam damage. Beam healing is just as bad for your image interpretation!

31.6.B Periodic Images

In discussing quantitative analysis, we have already noted how we can use the computer to identify similar features and combine them in order to reduce the noise. This technique has many possible variations. Again biological applications are leading the way with 3D crystallographic reconstruction, as illustrated by Downing and by Dorset and even correcting for distortions in the specimen, as shown by Saxton et al.

31.6.C Correcting Drift

Although drift is not as limiting on new machines, many older TEMs are still in use. Drift can be corrected now if the rate and direction of movement are constant. The computer can calculate the relative translation of two images and change the current in the image translation coils appropriately (which avoids moving the specimen). The difficulty is that the drift may not be linear. When implemented, such routines are already particularly valuable for frame averaging using a video camera. There will also be many applications for diffraction-contrast imaging as well as for X-ray or EELS analysis.

31.6.D Reconstructing the Phase

Although we are studying phase contrast, the image intensity doesn't directly give us phase information. Kirkland et al. showed as early as 1982 that the phase can be reconstructed by processing a defocus series. In their approach they use an iterative non-linear image-processing technique to reconstruct the complex

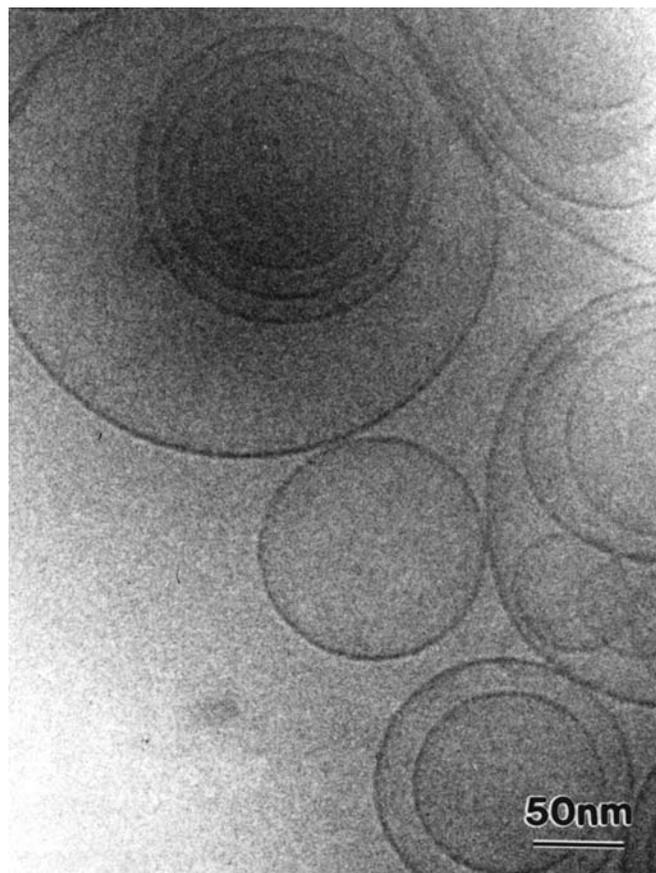


FIGURE 31.9. The image is from a highly beam-sensitive solution of surfactants in water. The solution has been frozen by plunging a film into liquid ethane and then transferring it to the TEM. The large circles are the surfactants that have aggregated to form vesicles; the concentration of the surfactants in the solution is just right for them to form lamellar structures (the circles seen flat on). Texture starts to appear in the image as soon as the beam interacts with the specimen due to crystalline ice formation and/or beam damage.

electron-transmission function. The technique was demonstrated using images of $\text{CuCl}_{16}\text{PC}$ (hexadeca-chlorophthalocyanine copper).

Five images from the experimental defocus series are shown in Figure 31.10A–E, together with the reconstructed transmission function plotted both as a real and imaginary part and then as an amplitude and phase part in Figure 31.10F–I. The projected structure of the known unit cell is also shown in Figure 31.10J. The phase image contains most of the structural information: it corresponds to the projected potential while the amplitude image contains features due to inelastic scattering. Notice in particular that we can now identify the benzene ring. This is one of the earliest published examples of the full-phase reconstructions. You must record such a series of defocus images if you want to do quantitative HRTEM. The companion text includes an extensive chapter giving details and an overview of how the process is practiced today.

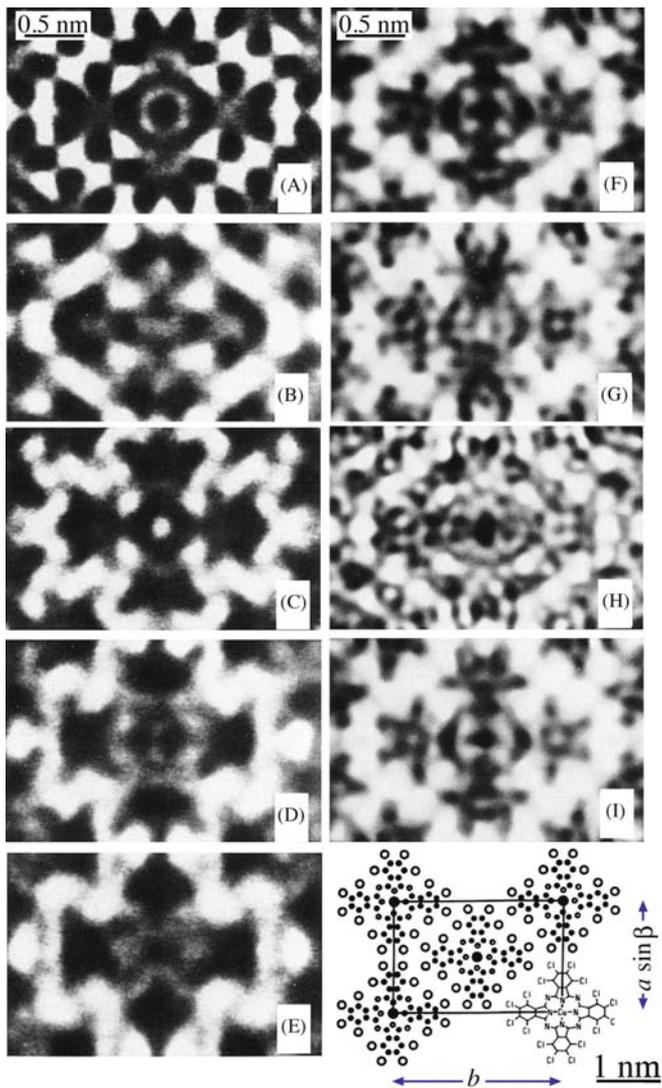


FIGURE 31.10. (A–E) Images from an experimental defocus series of CuCl_6PC ; the reconstructed transmission function plotted as both a real and imaginary part (F,G) and then as an amplitude and phase part (H,I); (J) the projected structure of the crystal.

31.6.E Diffraction Patterns

We have generally ignored the intensities in DPs because they are so strongly influenced by dynamical scattering. However, if the specimen is very thin, we can use the intensity in the SADP to carry out electron crystallography in the same way as in classic X-ray crystallography. As you can appreciate from Figure 31.11, particularly if the unit cell is large and the specimen examined is thin there is a great deal of information in the SADP but you can't get it all in one exposure. Hovmöller's group (see, e.g., Zou and the ELD program in Section 1.5) has provided a routine for analyzing such patterns and getting structure-factor information. For example

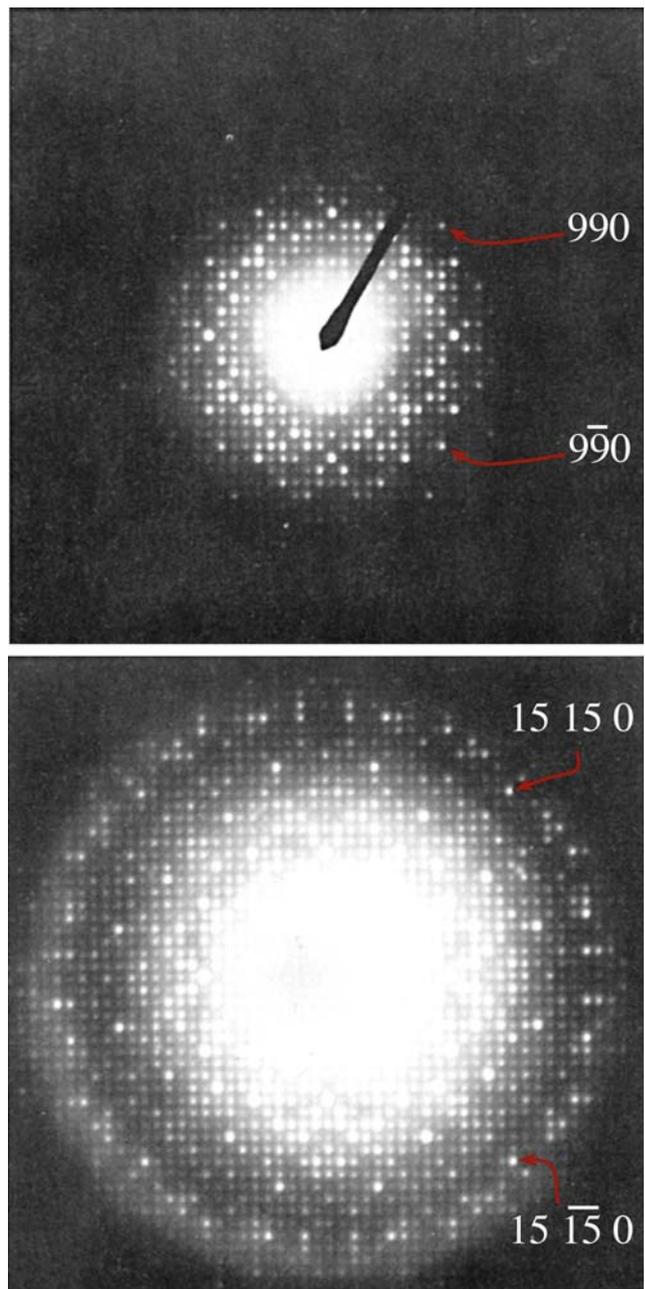


FIGURE 31.11. SADPs from $\text{K}_2\text{O} \cdot 7\text{Nb}_2\text{O}_5$ recorded using two different exposures. More than one exposure is needed to get all the information in the DP. The space group is $P4bm$ and $a = b = 2.75 \text{ nm}$. The $(15, 15, 0)$ reflections correspond to spacing of 0.13 nm .

- Several patterns are recorded using exposure times of 0.5 s to $>15 \text{ s}$.
- The patterns are digitized directly from the negatives using a CCD camera and a light-box for backlighting.
- The intensity of each film is calibrated using a calibration strip with 20 equal exposure steps.
- The intensity is measured for all the spots and the processing begins.

Thus far it is just like the classic analysis of XRD intensities. This digitization process is particularly

demanding, because each reflection typically covers an area of <0.5 mm diameter on the photographic film. You will need to be able to index three strong, but clear, reflections. The computer can then perform a series of functions

- Optimize the location of these points using a center-of-gravity approach, locate the origin, and index the rest of the pattern.
- Extract the intensities of each peak, taking care not to be misled by any shape effects of the specimen.
- Use reflections which are present on two successive negatives (since the intensities are now in digital form) to calibrate films recorded with different exposures and thus develop a very large dynamic range.

A cooled slow-scan CCD camera will give you a large dynamic range and better linearity than its room-temperature counterpart and should simplify this type of analysis. There are other complications in using electrons rather than X-rays for this kind of crystallography. While the Ewald sphere is still curved, as with X-rays, electrons can easily damage your specimen. However, the technique clearly has potential! Like all TEM techniques it can be applied to much smaller regions of the specimen than is possible for X-ray beams. We can also use the symmetry present in the SADPs. Because the specimen is very thin, this technique could be described as ‘kinematical’ crystallography and complements the ‘dynamical’ electron crystallography that we described for CBED patterns from thicker specimens in Chapter 21. The process of extracting intensities from DPs can now be carried out using the ELD software package (Section 1.5).

USING THE DP

The structure deduced from the HRTEM approach should generate the experimental SADP so it should be possible to use these diffraction data to further refine the structure.

If we can use the quantitative information available in DPs, we could combine this information with our experimental and simulated HRTEM images. The quantitative analysis of the DP is known as structure-factor-modulus restoration or reconstruction (Tang et al.). One limitation of this approach is that the specimen be sufficiently thin that diffraction is kinematical. Of course, this requirement is necessarily similar to the HRTEM requirement of the WPOA.

31.6.F Tilted-Beam Series

Having gone to great trouble to remove any beam tilt, we will mention how beam tilt can be used to extend the resolution of your microscope! The basic idea goes back to the tilted-beam lattice-fringe imaging we discussed in Section 23.3. Now you use a computer to combine information in different tilted-beam images. The method proposed by Kirkland et al. assumes that you know when the beam tilt is zero. You tilt the beam through different angles in well-defined directions so that you transfer information in overlapping regions of reciprocal space, as shown in Figure 31.12A; you also need the on-axis image, as shown in the tableau in

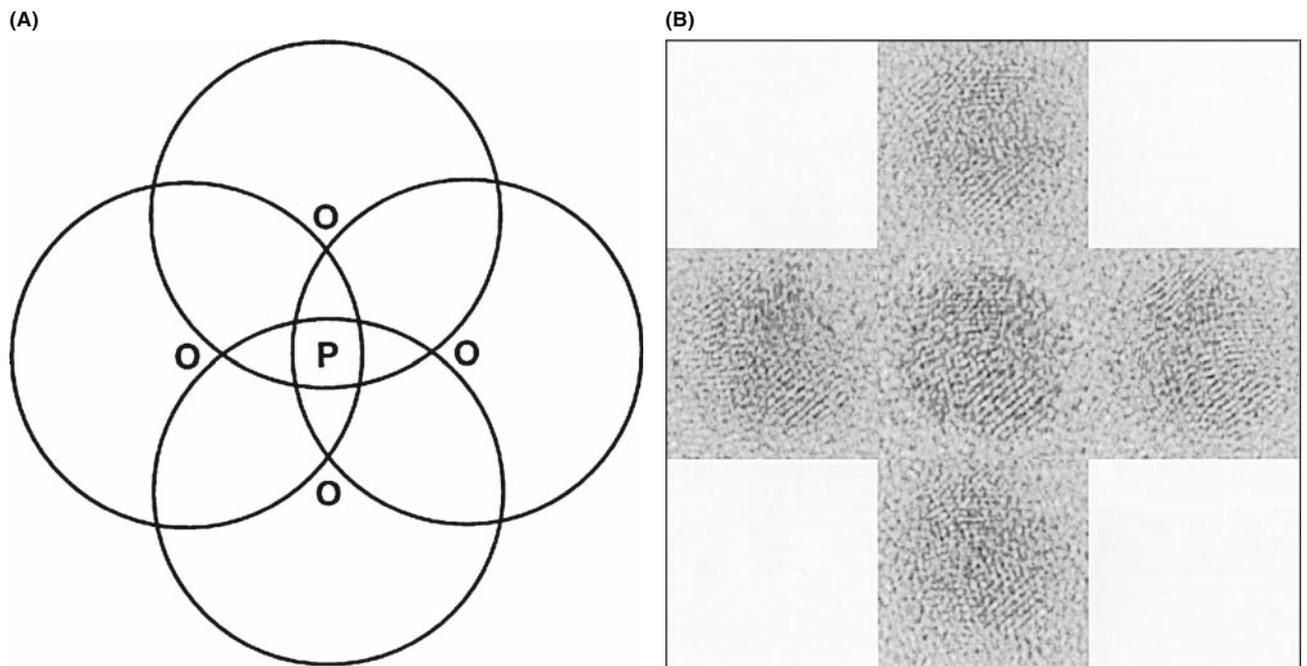


FIGURE 31.12. A method for extending the resolution of your TEM. Set the beam tilt to zero, then tilt the beam through different angles. (A) The four regions of Fourier space are shown by the four circles; O is each position of the tilted beam, P is the optic axis, and PO corresponds to the angle of tilt. (B) The five images used in the restoration arranged according to the beam tilt used in (A) with the on-axis image at the center.

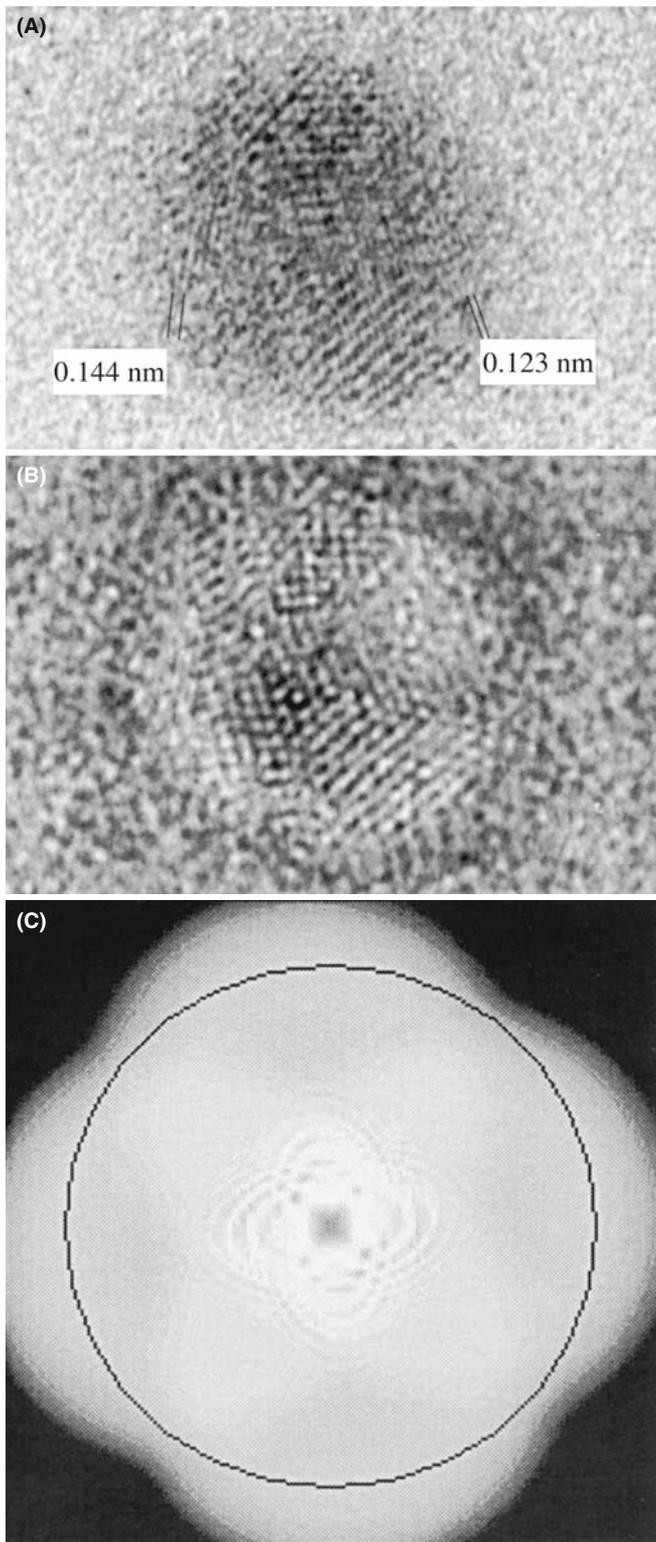


FIGURE 31.13. The restored image of a gold particle on amorphous Ge: (A) the amplitude (modulus) image showing 0.123-nm fringes, (B) the corresponding phase image, (C) the transfer function after restoration, plotted in two dimensions. The circle corresponds to 0.125-nm detail successfully transferred to the image. For such thin specimens, atom-dense positions have a reduced modulus [black in (A)] and an increased phase [white in (B)].

Figure 31.12B. Since it is important that the same area is imaged, a sixth (on-axis) image is recorded and correlated with the first (on-axis) image to check for drift and specimen degradation. You now need to restore the modulus and phase to create a higher-resolution image.

Kirkland's paper is a beautiful demonstration of the care needed in image processing. Even aligning the images is not trivial. However, the resulting restoration shown in Figure 31.13, demonstrates the potential of the technique: detail is present in the image at a resolution of 0.123 nm using a 400-kV microscope.

31.7 AUTOMATED ALIGNMENT

In the not-too-distant future, all TEMs will have automatic beam alignment, astigmatism correction, and readout for Δf . What makes this possible is the diffractogram analysis, a slow-scan CCD camera to digitize the image, and computer control of all the microscope functions. By microscope functions, we mean all lens currents, deflector currents, specimen drive, and aperture drives. The slow-scan camera is needed because the computer needs to make measurements on more than one ring in the diffractogram. So you will not actually have to sit in front of the microscope once you have loaded the specimen.

THE ADVANTAGE OF REMOTE TEM

The big advantage in remote control will not be that you can sit in Huntsville and operate a microscope in California but that you can locate the microscope in its controlled environment with no one opening the door to check if it is working (it was) or entering the room and thus changing the heat-load. If your specimen is not ideal or the TEM breaks down, you won't have to go for a walk on the beach in California but can continue word processing in Alabama.

Autoalignment is now available to everyone with a digital camera and the plug-in for Digital Micrograph (see Section 1.6). The early developments are summarized in the references. Your role is to select a suitable region of the specimen close to the area of interest; the area that you are interested in should ideally only be examined at low magnification. You will make the initial alignment manually and then ideally turn the process over to the computer. The computer will then adjust the astigmatism and correct the beam tilt independently and quickly.

Figure 31.14 shows how well and quickly this procedure can now be done. The different diffractograms in each tableau correspond to incremental changes in the beam tilt of 6 mrad in the x and y directions. The computer showed that the initial tilt error was 4 mrad, which

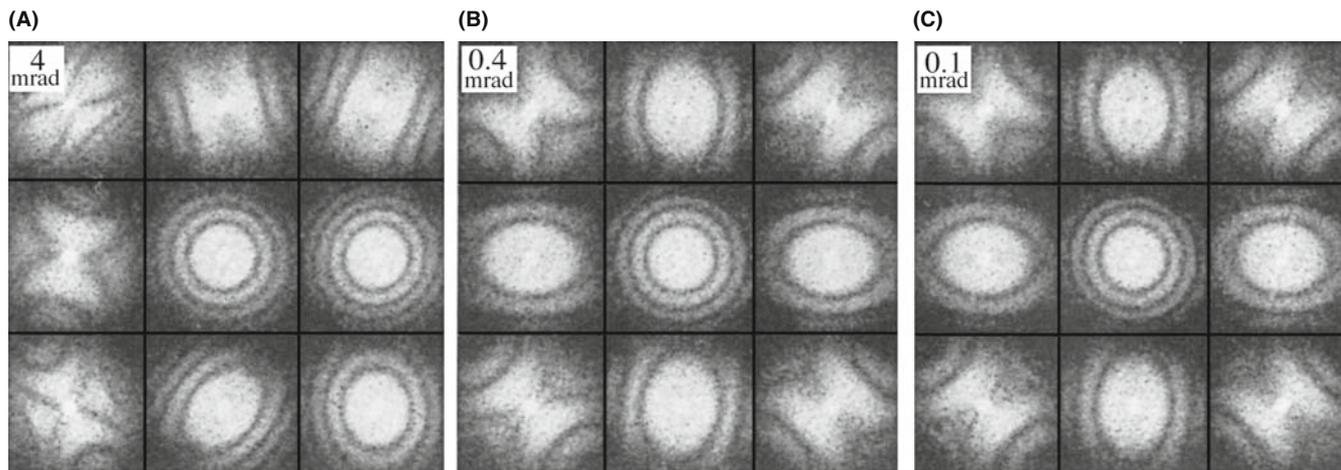


FIGURE 31.14. Using the computer to correct the beam tilt. (A) is the initial setting and (C) is the final. The incremental changes in the beam tilt are 6 mrad in the x and y directions away from the initial beam tilt in the (B). Note the central diffractogram is almost unchanged, emphasizing the need for computer-controlled tilting to give correct alignment. The misalignment in each setting is given in the upper left corner. (A) is the initial setting and (C) is the final.

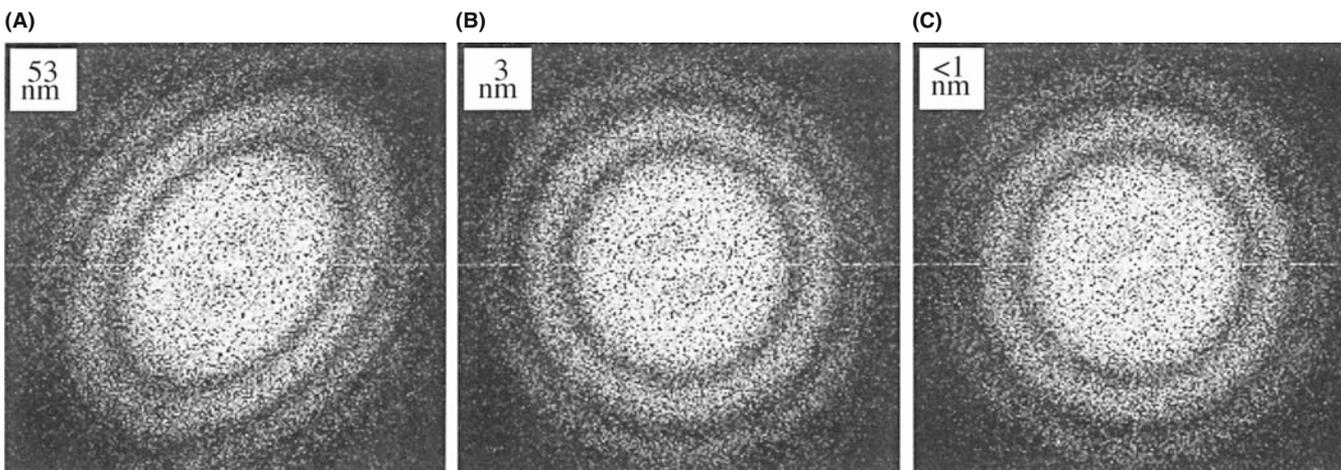


FIGURE 31.15. Diffraction patterns showing the astigmatism corrections made by the computer following a similar procedure to that shown in Figure 31.13. (A) is the initial setting and (C) is the final (and is not the same as (B)). (C) shows that the HRTEM is now well stigmatized. The misalignment in each setting is given in the upper left corner.

was reduced to 0.4 mrad after one pass and <0.1 mrad after the second pass. Each pass took only 28 s! The astigmatism shown in Figure 31.15 was initially 53 nm. It was reduced to 3 nm after one pass and to <1 nm after the second pass. For this correction, each pass took only 8 s. Even the most experienced operator can't match this speed or accuracy for either correction, and both corrections are now *quantitative*.

The defocus value is then found by calibrating the image with minimum contrast occurring at Δf_{MC} . The value of Δf_{Sch} can then be found when the image contrast is a maximum. Although the method described here uses the diffractogram, a corresponding approach can be followed by analyzing variations in the contrast of the image. This technique has been described by Saxton et al. and uses a method of cross-correlating

pairs of images recorded at each focus setting of the microscope. The reason for cross-correlating images is to remove the effects of electron shot noise; variations due to the photographic emulsion are avoided by using the slow-scan camera.

31.8 QUANTITATIVE METHODS OF IMAGE ANALYSIS

In the next six sections, we will go through several special topics and illustrations of image processing in HRTEM.

- Pattern recognition
- Parameterizing the image
- Chemical information from images

- Measuring the fit
- Quantitative comparison of simulation and experiment
- Fourier techniques

Our discussion will draw heavily on the work of a few pioneers in this field; we will also emphasize that, although this subject is still in its infancy, it is developing rapidly. The main cause for the delay in its application in materials science has been the lack of affordable fast computers and the feeling that everyone must write their own image-processing program; the latter is not true and is certainly not recommended. At this time, we can summarize the situation as follows

- Quantitative analysis is difficult, often tedious, and invariably time consuming.
- You must understand the basic ideas of image theory before you start processing images.
- Your analysis is only as good as your image and your image is only as good as your specimen.

We gave some information on software in Section 1.6 and expand on this in Section 31.16.

31.9 PATTERN RECOGNITION IN HRTEM

The most obvious feature of most HRTEM images is that we see patterns of white, gray, and black dots or other shapes. If the pattern is perfect everywhere, your specimen is probably a single crystal with no defects, no thickness variations, no variation in atomic composition, and no use. If it is not perfect, then we can use pattern recognition to quantify the variations.

PATTERN RECOGNITION

The idea is to take/make a template, move it across your image, and measure how closely the image resembles your template.

Clearly you need a computer for this! Your template needs to match the magnification and rotation of the pattern you are examining. Then you need a method to say how close your match is, i.e., you need to know your 'goodness of fit.' We will go through some basics here, but strongly recommend that you consult the list of original papers given at the end of this chapter when you are ready to apply this technique.

We can illustrate the approach following Paciornik et al. The large rectangle shown in Figure 31.16 represents your digitized image and could be $1\text{ k} \times 1\text{ k}$; remember the numbers indicate pixels. The small rectangle represents your template. This template might be a small area of the pattern or a simulated image, in which case it might be a 128×128 pixel template. If the template is taken from your image, then you have

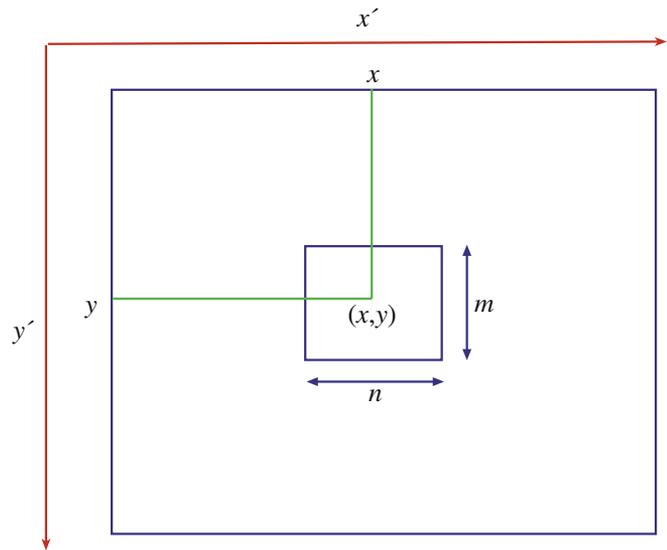


FIGURE 31.16. The large rectangle represents the digitized image, size $(x' \times y')$; the small rectangle, $(m \times n)$ pixels, represents the template used in the cross-correlation calculation. The small rectangle is moved to different (x,y) positions during the process.

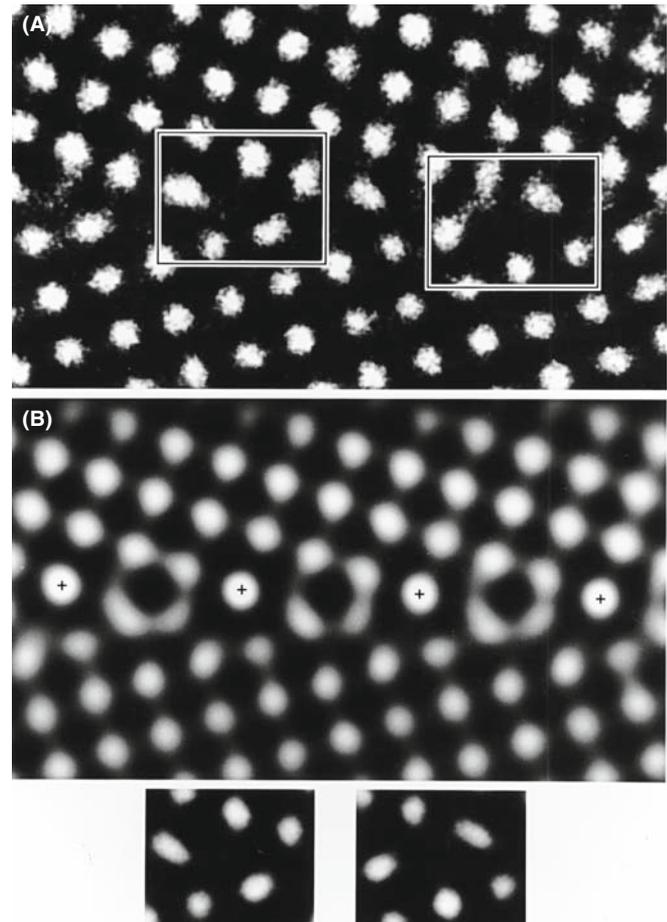


FIGURE 31.17. Analysis of a small region of a $\Sigma = 5$ tilt boundary in TiO_2 . The two small boxed regions in (A) are only present at the boundary; these are used as the templates for the cross-correlation method. (B) The cross-correlated image. The small rectangles at the bottom of the figure are the low-noise-averaged images of the GB templates.

already got the right magnification and rotation. If not, you have to set these first, but remember that the lens will still cause distortions. We will return to this problem shortly. Klug actually used pattern recognition in the work that led to his Nobel Prize.

REAL-SPACE APPROACH

We only look at the image.

This is a real-space approach. The process is best understood by an example. Figure 31.17A shows an HRTEM image of a small region of a $\Sigma = 5$ tilt boundary in TiO_2 . The two small boxed regions appear only at the boundary and are selected as templates. The matching process has then been carried out and the new image is shown in Figure 31.17B. Having found all the regions which match the template, we could then take the average of these to produce low-noise images of the GB templates. The final step is the comparison of these templates with models of the GB structure. There are two important points to remember

- When you average images, you implicitly assume that all the images are the same except for random noise.
- Don't forget our discussion in Chapter 30 of interface grooving and the problems associated with interfacial segregation.

31.10 PARAMETERIZING THE IMAGE USING QUANTITEM

In general, the thickness or chemistry will vary as you cross the specimen, i.e., the projected potential varies across the specimen. This means that one template will only match a small area, so you have to use many templates. These templates could, in principle, be totally empirical, but to be quantitative we must derive them from image simulations. This approach has been described for two special cases in papers by Kisielowski et al. and Ourmazd et al.

31.10.A The Example of a Specimen with Uniform Composition

In QUANTITEM, the results of Chapter 30 are summarized by a general equation linking the intensity and all the imaging (S_i) and materials (P) parameters

$$I(x, y) = F(P(x, y), S_i) \quad (31.5)$$

This equation just tells us that the intensity depends on the imaging conditions and on the specimen. For a particular set of imaging conditions, S_i will be known (more or less) and we'll call it S_i^0 . Then we can write that

$$I(x, y) = F(P(x, y), S_i^0) = F^0(P(x, y)) \quad (31.6)$$

The basis of this approach is quite straightforward, potentially powerful but almost never used

- Define the function F^0 for each image that you may obtain.
- Then construct a set of templates for your matching process.

Providing you stay within one extinction band, F^0 will be directly related to the projected potential of the specimen. A helpful simple analogy for this process is illustrated in Figure 31.18. The function F^0 describes the path of a swinging pendulum as it varies with time (Figure 31.18A). Each value of F^0 corresponds to a snapshot of the pendulum, so if you plot F^0 you can 'see' the path of the pendulum (Figure 31.18B). The velocity of the pendulum is related to the density of points along the path. So it should be possible to plot out the function F^0 from a single lattice image even if you don't know the microscope parameters used to form the image.

Yes, there are limitations and conditions and we'll discuss them later. All we need now is a method for representing each image by a snapshot of the pendulum: we have to *parameterize* the image. This process is the key to the technique. Manipulating and quantifying, in principle, thousands of images, each requiring 4 M byte of memory is not a fast process, even if you do have that much memory. If we could characterize each image by a few numbers (a vector or parameter) the comparison process could be much faster.

We separate the image into unit cells and digitized these to give many templates which are n pixels by m pixels, as shown here for one cell (Figure 31.18C–F). If we define N to be $n \times m$, then we have N numbers for the N pixels, where each number represents a gray level. Now the N numbers are regarded as the N components of an N -dimensional vector. (The math is not complicated but don't try to visualize this vector.) So now all the information in each unit cell is represented by a vector in N -dimensional space. The function F^0 describes how these N -dimensional vectors (the path) change as the projected potential changes.

PARAMETERIZING IMAGES

This may remind you of using the Hough transform to quantify EBSD patterns from the SEM.

The next step is to define a reference frame for these vectors. Three basis vectors are derived from the experimental image. Ourmazd et al. argue that three basis vectors will be sufficient, as we can show in the

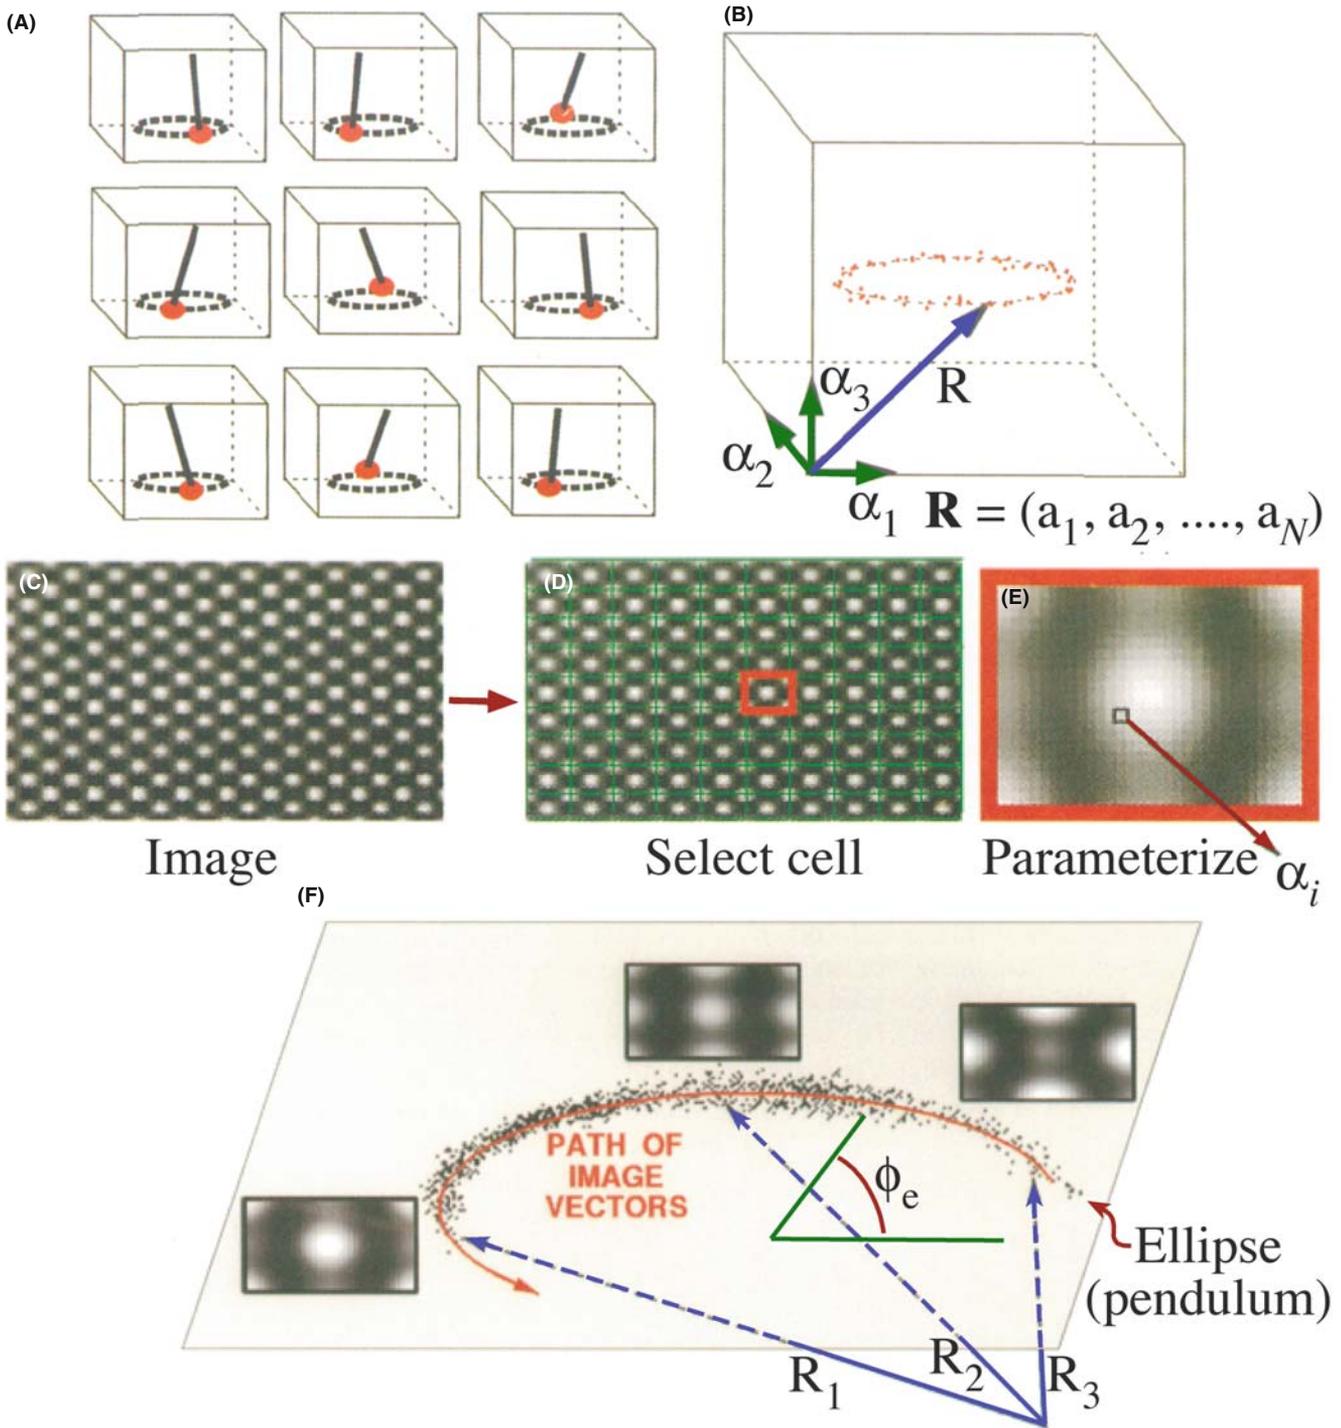


FIGURE 31.18. The principle of vector parameterization used in QUANTITEM. Each HRTEM image is represented by a single vector \mathbf{R} which has N dimensions. (A) and (B) show the ‘swinging pendulum’; (B) shows its ‘path’; (C–E) the image is separated into unit cells and digitized to give $(n \times m)$ pixel templates; (F) three vector-parameterized images (\mathbf{R}_1 , \mathbf{R}_2 , \mathbf{R}_3) of a wedge-shaped specimen of Si at different thicknesses.

following way. We will be using a low-index zone axis for any HRTEM analysis. Then we have three types of images

- The background, \mathbf{R}^B , due to the direct beam, O .
- A single-period image, \mathbf{R}^S , due to the interference between O and the strongest reflections, G_i .

- A double-period image, \mathbf{R}^D , due to the interference between these strong G_i reflections.

Each of these \mathbf{R} terms is a vector which represents an image. Any image we can form must be a combination of these three types of image, so a general image, G , can be written as

$$\mathbf{R}^G = a_G \mathbf{R}^B + b_G \mathbf{R}^S + c_G \mathbf{R}^D \quad (31.7)$$

Each of the basis vectors (images) can be expressed in the same manner

$$\mathbf{R}_i^T = a_i \mathbf{R}^B + b_i \mathbf{R}^S + c_i \mathbf{R}^D \quad (31.8)$$

giving three vectors for $i = 1, 2,$ and $3.$

We can, in principle, turn these equations around (invert them) to define any vector \mathbf{R}^G in terms of the basis vectors

$$\mathbf{R}^G = \alpha_G \mathbf{R}_1^T + \beta_G \mathbf{R}_2^T + \gamma_G \mathbf{R}_3^T \quad (31.9)$$

which is what we wanted to show.

Ourmazd et al. point out that this treatment gives three important results

- The vector notation allows us to parameterize the lattice image.
- Projecting the vectors onto planes and/or paths aids noise reduction.
- Any noise which remains can be quantified.

The result of vector-parameterizing an experimental image of a wedge-shaped specimen of Si is shown by the three images corresponding to $\mathbf{R}_1, \mathbf{R}_2,$ and \mathbf{R}_3 in Figure 31.18F.

31.10.B Calibrating the Path of R

In order to relate any image to the projected potential, we have to calibrate the curve showing the path of $\mathbf{R}^G.$ This is where the image simulation comes in. We start with the vector-parameterized analysis of the series of simulated images from the Si wedge. Each point on the curve corresponds to an image of the unit cell, and thus to a vector $\mathbf{R}^G.$ The ellipse has been fitted empirically, and the thickness of the cell has been increased by 0.38 nm for successive calculations. The points are closer together in some parts of the plot because, as we saw in Chapter 30, some characteristic images appear for a wider range of thicknesses. Now we have a way to quantify this ‘experimental’ observation. What the ellipse does is to allow us to parameterize the path in terms of the phase angle of the ellipse $\phi_e,$ shown by the ellipse in Figure 31.18F. Thinking back to the pendulum analogy, the path parameters are the image version of the coordinates for the harmonic oscillator.

Now the parameters of the curve (ϕ_e) can be obtained from a series of images. We can vary the material and in each case examine three other variables

- The orientation of the specimen (i.e., the zone axis).
- The defocus, $\Delta f,$ of the objective lens.
- The specimen thickness.

The remarkable result is that when we plot ϕ_e versus the thickness, normalized by the extinction distance, we obtain a straight line. The explanation for this result is related to the fact that only a small number of Bloch waves usually contribute to the image, as we saw in Section 30.6. In materials such as YBCO, this is not the case, and the curve is not a straight line.

This analysis gives more detail than most students will ever need, but it introduces a concept that could (should) be more widely recognized: the power of parameterizing the image.

31.10.C Noise Analysis

Noise in the image moves the vector off the ellipse; so, we have a method for analyzing the noise. If the noise moves the vector exactly along the ellipse, we can’t analyze the noise, but that is quite unlikely since the noise would then be accurately mimicking a change in projected potential. So this parameterization gives a way to reduce the noise by a factor of \sqrt{N} which for a 10 pixel \times 10 pixel cell is a factor of 10!

The analysis given by Ourmazd et al. then shows that, in the case where only two Bloch waves are excited, the image intensity, $I,$ can be expressed as

$$I = B + S + D \quad (31.10)$$

where $B, S,$ and D are the contributions from the background, single interaction, and double interaction as we defined them above. The point (B, S, D) does indeed describe an ellipse which lies on a plane independent of $\Delta f.$

The value of this approach can be appreciated if you look at the examples shown in Figure 31.19. In the first example (Figure 31.19A and B), the technique has been used to provide a map of the roughness of the Si surface. The experimental image looks really uniform until you analyze it using this method, when you can discern the roughness at the 0.5-nm level; remember that there are two surfaces.

As you know from earlier discussions, changes in chemistry produce effects which are similar to changes in thickness, because they change the projected potential. In terms of the present analysis, the effects are different: composition changes cause changes in the ellipse and in ξ (notice that there is no subscript, since this ξ is a many-beam value).

The method is more limited, in this respect, to change in composition but can be used if the thickness and roughness are known, i.e., if you can measure the roughness elsewhere on your specimen (using a known reference cell) and infer it for the area you want to analyze. (Warning lights should be flashing.) The approach is as follows

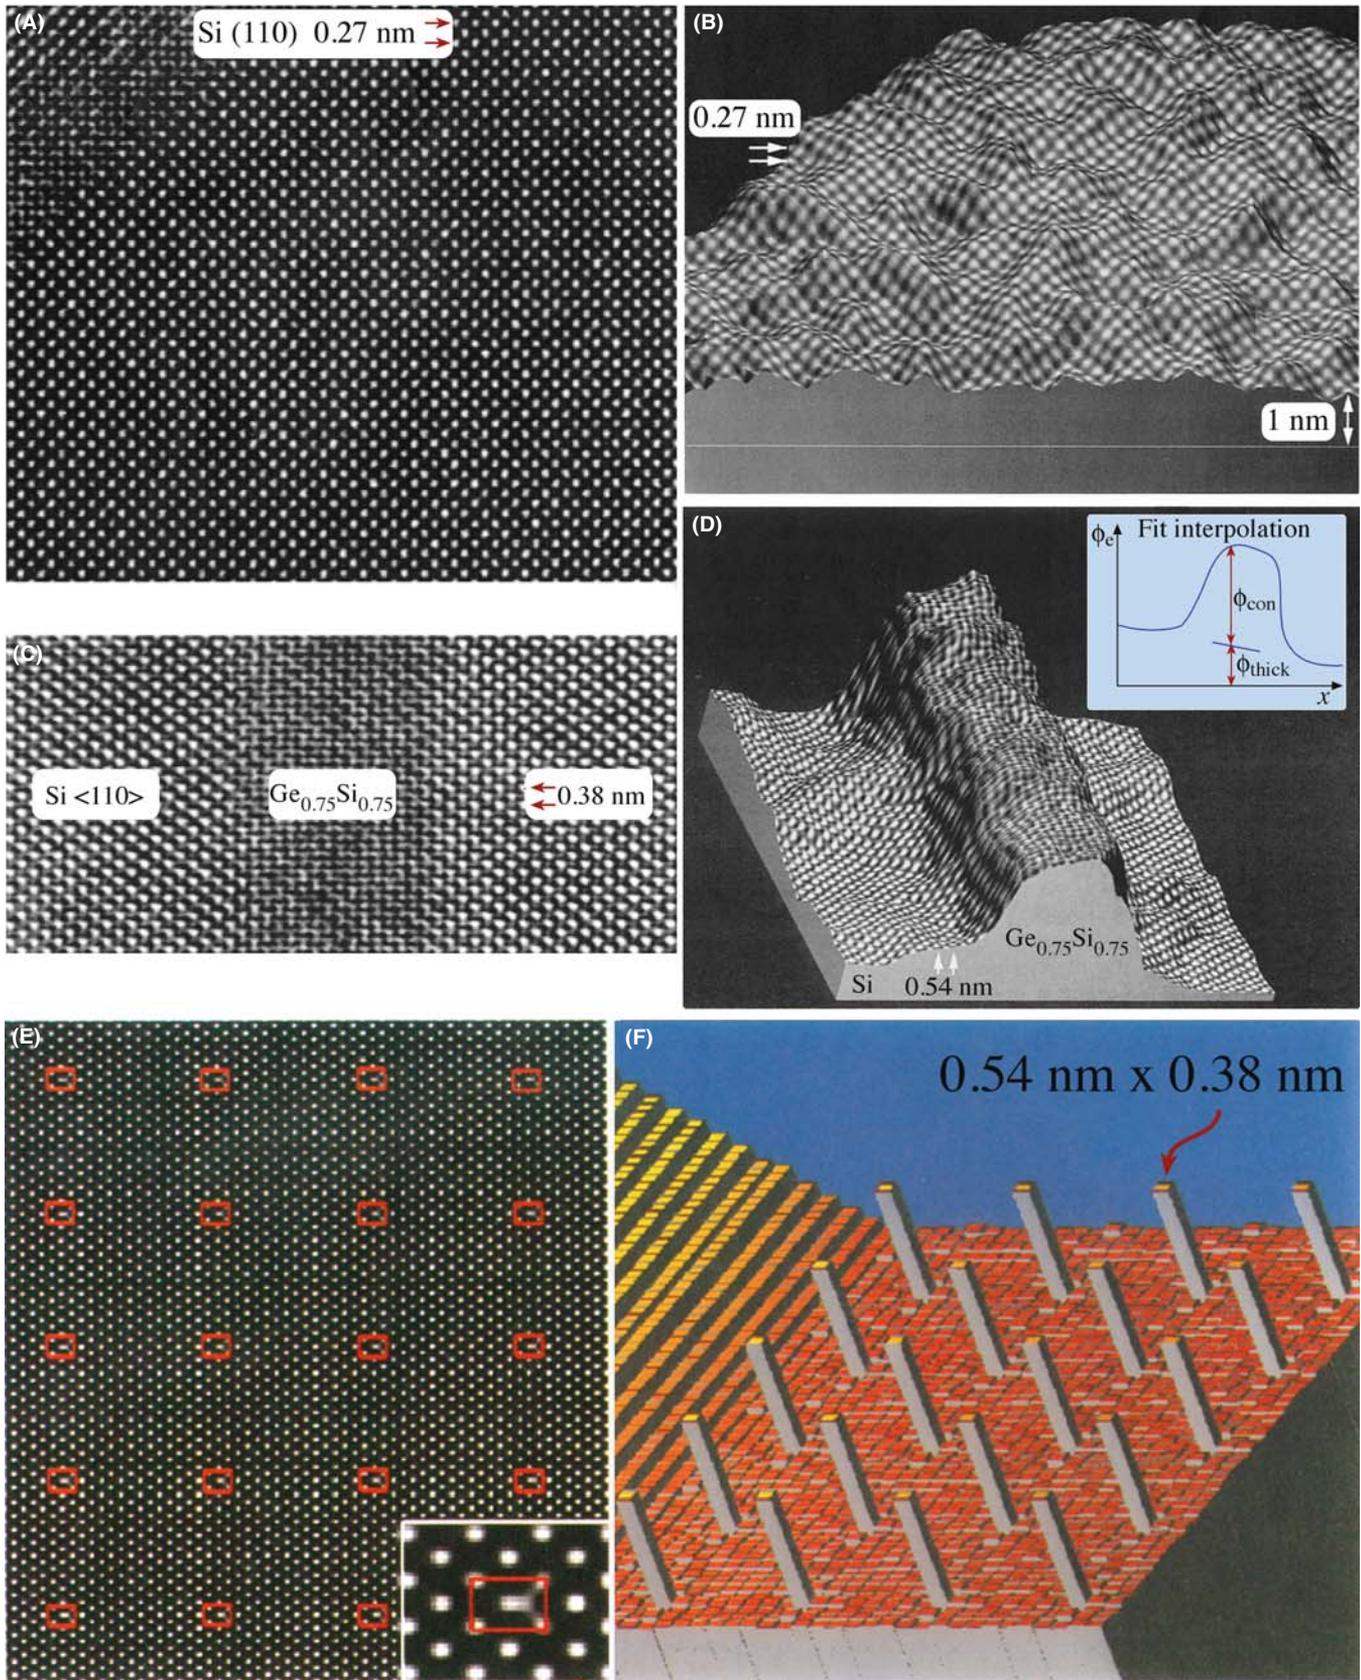


FIGURE 31.19. Examples of the application of QUANTITEM: on the left is the image, on the right is the QUANTITEM image. (A, B) Mapping of the roughness of the Si surface covered by SiO_2 ; (C, D) a layer of $\text{Ge}_x\text{Si}_{1-x}$ in a matrix of Si, the inset shows the plot of ϕ_c versus x ; (E, F) analyzing a simulated image of columns of Ge (a δ -function in concentration) in Si.

- Use QUANTITEM to measure the advance in ϕ_e at your target cell relative to your reference cell.
- Subtract $\Delta\phi_e$, which is due to a thickness change.
- Then the rest of the change in ϕ_e must be due to changes in ξ . If you know how ξ varies for different compositions, you have determined the local composition.

The second example shown in Figure 31.19C and D is a near-perfect application for the technique, since the elements in the alloy $\text{Ge}_x\text{Si}_{1-x}$ are randomly located on the lattice sites. The slope tells us how abruptly the composition varies.

NOISE

You can test the potential resolution of the technique, its sensitivity to the alignment of the beam, bending of specimen, and beam divergence in the usual way by creating model structures, simulating the images, and then analyzing them.

The third example in Figure 31.19E and F shows that the potential resolution is superb (we can see the quantum dots) but beam tilt can cause 10% errors in thickness measurement. The conclusion is clear: as always, you only get the best results if your specimen is ideal and both the microscope and the specimen are perfectly aligned. Note, however, that the technique has not yet been successfully applied to a wide range of materials, but it is complementary in many ways to STEM Z-contrast (see Figure 22.15).

31.11 QUANTITATIVE CHEMICAL LATTICE IMAGING

This technique is an instructive tool but is perhaps less necessary now that STEM is commonly available. The technique has not been widely used since the software was not readily available.

This technique uses the approach described in Section 30.4, but can only be applied to materials where we have chemically sensitive reflections, which we discussed in Section 16.4. We used these reflections in Section 29.4 to produce chemically sensitive DF images. In HRTEM, the chemically sensitive reflections not only contribute to the overall image but they will generally have a different dependence on thickness, too.

This effect is shown in Figure 31.20A–D for AlAs and GaAs, which have identical structures: the 002 reflection is allowed for both, but is stronger for AlAs since F , the structure factor, is proportional to $f_{\text{III}} - f_{\text{V}}$. You can see that, under the conditions chosen for this comparison, the intensity of the 022 reflections is also

very different for the different thicknesses. The sort of image we can analyze with this approach is illustrated here by the quantum wire. We want to know how abruptly the composition changes at the interface. In this example, the ideal GaAs and $\text{Al}_{0.4}\text{Ga}_{0.6}\text{As}$ unit-cell images are first characterized by the two vectors, \mathbf{R}_{GaAs} and $\mathbf{R}_{\text{AlGaAs}}$, following the approach described in Section 31.10. In this case the cells were simulated, divided into 30×30 pixel arrays (so $N=900$), and then \mathbf{R} was plotted. The information content is contained in θ_C . As before, we can directly assess the noise in such an image. So how is the direction of \mathbf{R} dependent on composition?

The technique is explained using the three known simulated templates (GaAs, AlGaAs, and AlAs); each one produces a vector \mathbf{R}^t . Although the vector for the intermediate composition does not lie in the plane, it can be projected onto this plane to give a unique vector for certain ranges of thickness. Since this is a complex procedure, you'll find the 'image flow chart' shown in Figure 31.21 helpful.

- The experimental image is first digitized; the image contained approximately 25×25 unit cells and used a 514×480 frame buffer.
- Next, the image must be separated into individual cells.
- The pair of templates shown is then used to calculate the angular positions of the \mathbf{R} vectors for all the unit cells. Such templates can be calculated or taken from known areas of the specimen.
- These \mathbf{R} vectors are characterized in terms of where they cut through a plane (D) (see Figure 31.21D also).

The maximum chemical difference determines how far apart the two principal distributions can be (see Figure 31.21E). Since the image is now fully parameterized we can do the statistics and finally invert the angular data to give the compositions.

This technique has enormous potential, but you must also remember that it is susceptible to all the drawbacks inherent in HRTEM. The advance is that now you can put numbers on those effects. The technique is material-specific, but if you know your material, you can combine image simulation and this processing method to examine what will be the limiting factors for your material. You can construct a test image like that shown in Figure 31.19E. If your specimen is ideal, you could, in principle, easily detect a column of Al in a mainly GaAs matrix without any 'spreading' due to the electron beam and in this respect alone makes QUANTITEM better than atomic-resolution X-ray mapping and equivalent to atomic-column EELS (see Part 4). This technique could be applied to other materials.

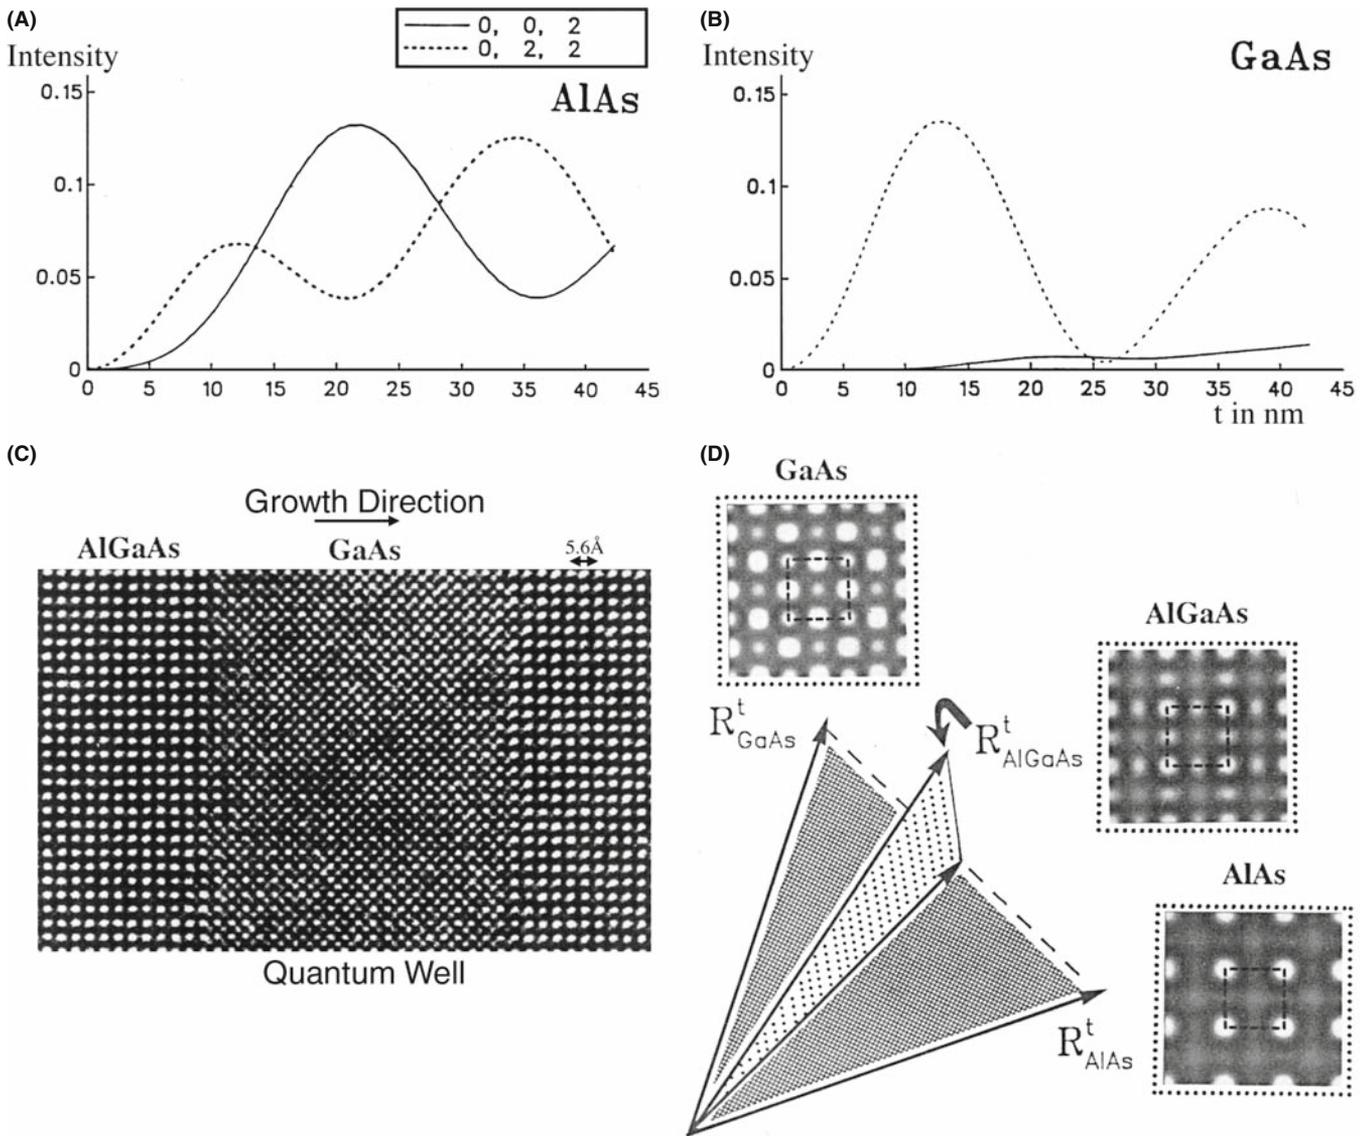


FIGURE 31.20. (A,B) Variation in intensity of the (002) and (022) beams along [100] in AlAs and GaAs (400 keV); (C) chemical lattice image of a layer of GaAs between two layers of $\text{Al}_x\text{Ga}_{1-x}\text{As}$ ($x = 0.4$); (D) templates simulated for different values of x each produce a vector \mathbf{R}^t .

31.12 METHODS OF MEASURING FIT

There are two methods presently used to obtain a measurement of how well the data and the model match. In statistics, we determine this ‘goodness of fit’ using cross correlation or a least-squares refinement (Section 31.13).

CROSS CORRELATION

The cross correlation gives a measure of how similar two images (or signals) are. The math process looks very much like a convolution of two functions but we don’t take the Fourier transform. What’s special is that the computer does it easily as long as our images are numbers!

LEAST-SQUARES

A method of fitting a curve to a series of values.

In this section we’ll use the cross-correlation method to compare an $n \times m$ pixel template (see Section 31.9) with every possible $n \times m$ rectangle in the image. The computer moves the template across the image one pixel column at a time, then shifts down one pixel row, and repeats the exercise. The cross-correlation function (CCF) gives the goodness of fit or a ‘measure of similarity’ between the template and each $n \times m$ image

$$CCF(x, y) = \frac{\sum_{x'} \sum_{y'} [i(x', y') - \langle i(x', y') \rangle] \cdot [t(x' - x, y' - y) - \langle t \rangle]}{\sqrt{\left\{ \sum_{x'} \sum_{y'} [i(x', y') - \langle i(x, y) \rangle]^2 \sum_{x'} \sum_{y'} [t(x' - x, y' - y) - \langle t \rangle]^2 \right\}}}$$

(31.11)

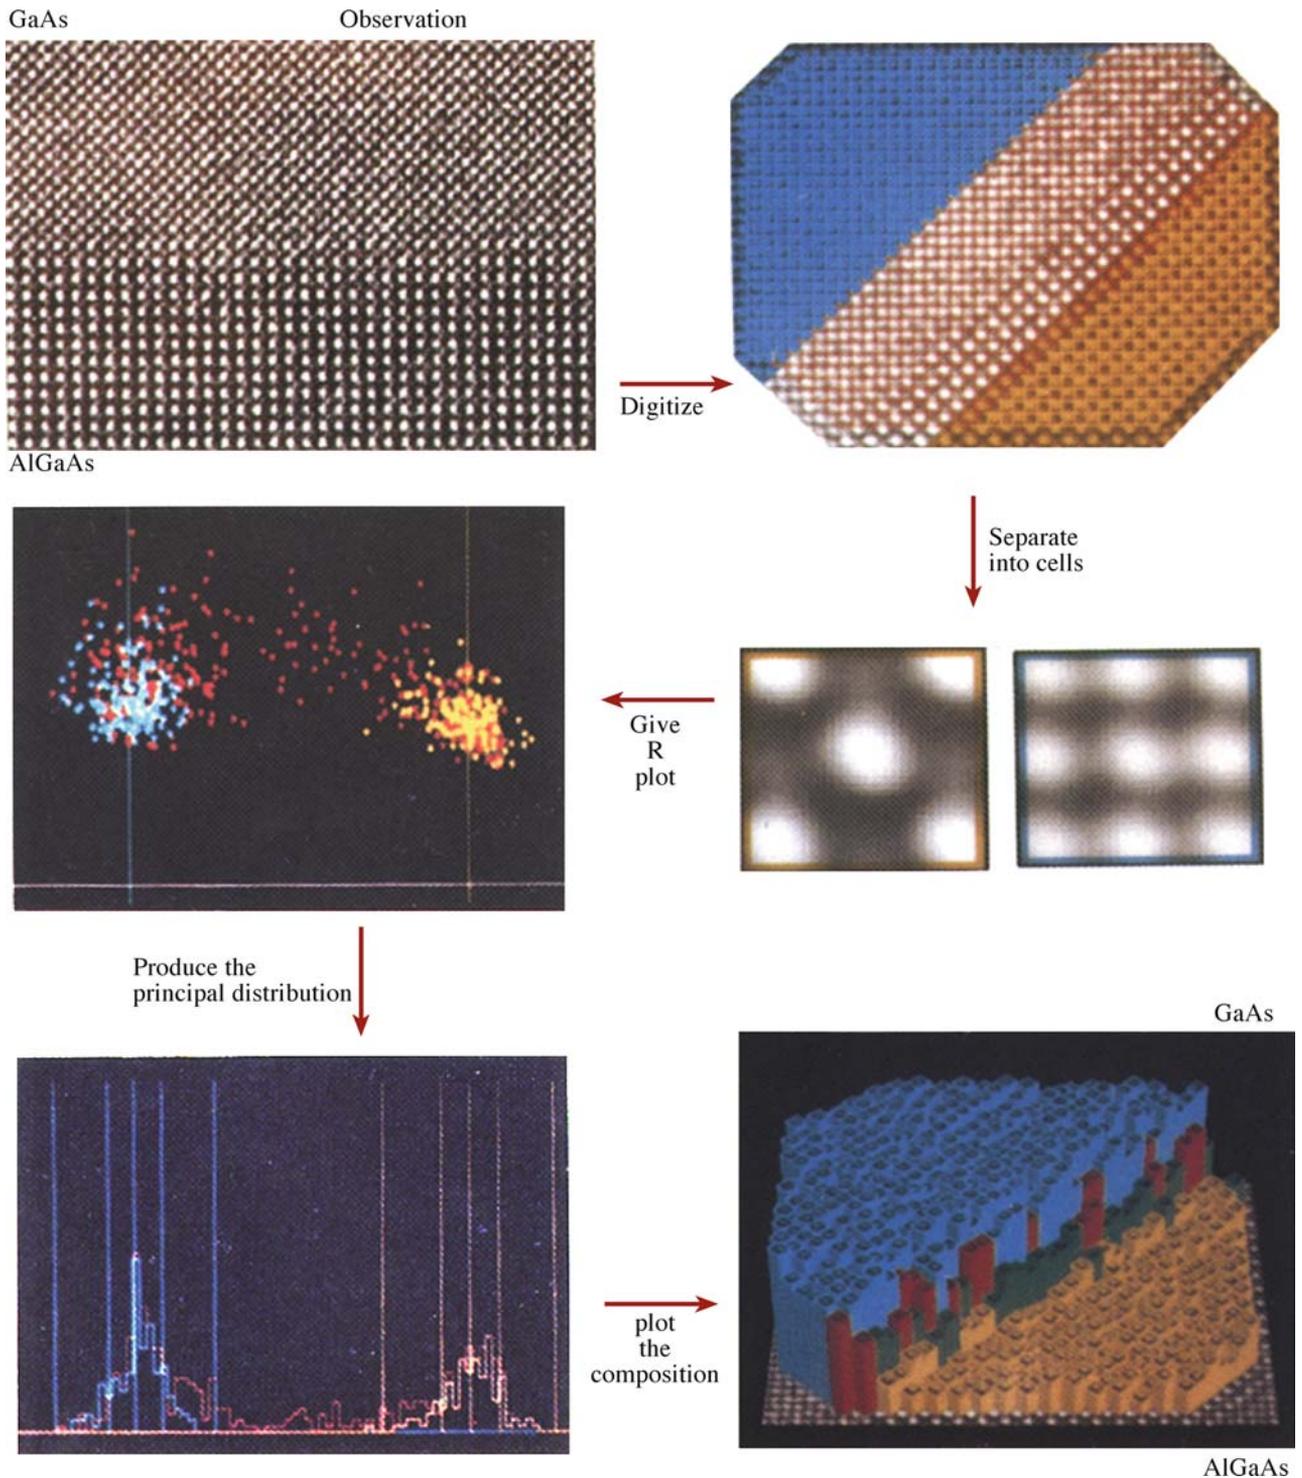


FIGURE 31.21. A flow chart summarizing the pattern-recognition procedure.

In this equation x varies from 0 to x_{\max} , y varies from 0 to y_{\max} .

- $i(x', y')$ represents the image.
- $t(x', y')$ represents the template.
- $\langle t \rangle$ is the average value of the pixels in $t(x', y')$; it is computed just once.

- $\langle i(x, y) \rangle$ is the average of $i(x', y')$ in the region coincident with the current location of t .

The summations are taken over the coordinates common to both i and t . The origin of the image is at its top left corner and the origin of the template is at its center. In this equation, the denominator is a

normalization factor, so that the CCF will not depend on any difference in the intensity scale between the template and the image.

We can rewrite equation 31.11 as the dot product of two vectors \mathbf{t} and \mathbf{i} , which gives us the $n \times m$ component of the template

$$\text{CCF}(x, y) = \cos(\theta) \frac{\mathbf{t} \cdot \mathbf{i}}{|\mathbf{t}| |\mathbf{i}|} \quad (31.12)$$

Now we can plot the CCF as a map of our image and then examine it to deduce where there is a particularly good match. Since the CCF value varies from 0 to 1, we can plot out the number of times each particular CCF value occurs, as illustrated in Figure 31.22. The two peaks in this curve correspond to the best fit and the worst fit, so the distance between them gives a measure of the 'discrimination signal.' From the width of the peaks we have a measure of the noise and hence a signal-to-noise ratio. The regions of good fit can be combined to produce a better template, and the process repeated, giving the dashed line. A second measure of the noise is then given by how far the good peak differs from unity. A particularly nice feature of this approach is that the procedure is available as a plug-in module for Digital Micrograph (see Section 1.6). The template could alternatively be a simulated image and the process repeated for a series of different thicknesses and/or defocus values. When you want to learn more about correlation techniques, see the article by Frank.

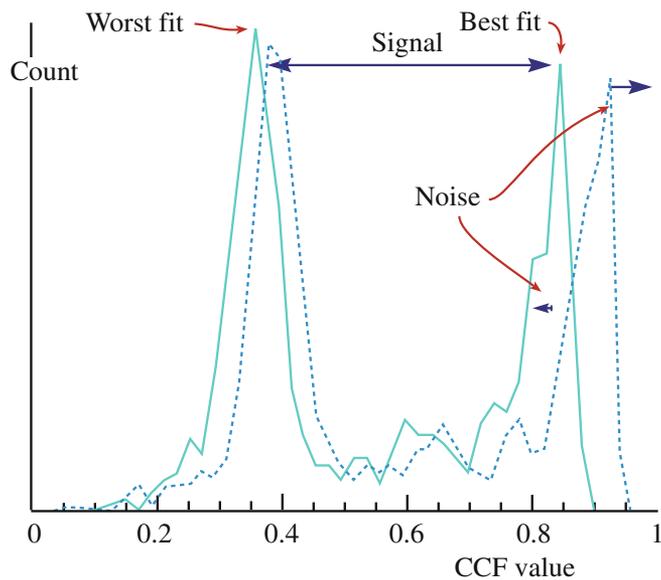


FIGURE 31.22. A plot of how often a particular CCF value occurs. The two peaks in the continuous curve are the best fit and the worst-fit: their separation gives a measure of the discrimination signal; the width of the peaks gives a measure of the noise and hence a signal-to-noise ratio. The plot can be redrawn after repeating the process (dashed line) to estimate the improvement in signal-to-noise ratio.

31.13 QUANTITATIVE COMPARISON OF SIMULATED AND EXPERIMENTAL HRTEM IMAGES

If we want to compare simulated and experiment images quantitatively, we really should modify our usual approaches to both simulation and experiment (King and Campbell 1993 and 1994). When doing the simulation, most programs automatically adjust the gray scale for each image so that darkest is 1 and brightest is 0 (or vice versa). This means that two simulated images might appear similar even though you would hardly see the pattern in one if both appeared on the same negative. In a similar way, we usually print an image to be as clear as possible using the full range contrast of the photographic paper.

We need methods for normalizing these procedures if we want to make quantitative comparisons. The solution for the simulation is simple. For the experimentalist, it means recording extra data while you're at the microscope. After recording the image, you record another image with the specimen removed. You then use this image to scale the lattice image such that you correct for variations in intensity across the field of view and the non-linearity of the response from the photographic film. Figure 31.23 illustrates the experimental transmittance for Kodak SO-163 film, 400-keV electrons, plotted against the digital value on a CCD array. Of course, you must process both images at the same time. This is called the 'flat-field' correction; a slow-scan CCD camera would simplify this procedure at the cost of reducing the area you examine.

When you analyze the image, you'll find out if the area you photographed was correctly aligned. Since your image takes only 2 seconds or so to record, you

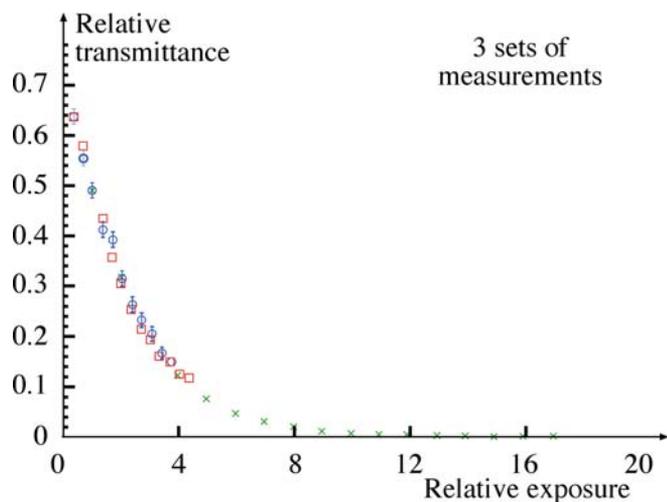


FIGURE 31.23. Plot of transmittance versus relative exposure measured using a CCD camera to digitize images from Kodak SO-163 film. (The symbols indicate three different microscopes.)

may risk several exposures using this technique. You are now comparing numbers, so you can use a least-squares fit where the residual $f_i(x)$ is defined as

$$f_i(x) = \frac{[f_i^{\text{obs}} - f_i^{\text{calc}}(x)]}{W_i} \quad (31.13)$$

and your task is to minimize $f_i(x)$. The difference between the intensity in the experimentally observed i th pixel and its calculated value will be zero if everything has been scaled correctly, the imaging conditions (Δf , C_s , etc.) are correct, and you have the right structure.

HRTEM AND LOW DOSES

Don't forget that since HRTEM uses higher voltages, the perfect image will only be recorded from an area of the specimen that has only seen the beam while you recorded the image! So you should always use low-dose techniques for quantitative imaging (but others don't).

Let's say W_i is the image which represents the error bars for pixel i . Then we can write that

$$W_i = \min \left[\sum_{i=1}^N f_i(x)^2 \right] \quad (31.14)$$

This equation defines the non-linear least-squares problem. We use x to summarize a set of parameters (Δf , C_s , the model, etc.); N is the number of pixels in the image. Fortunately, this analysis is now routine statistics. You'll need a computer program to tell you how good the first guess was, make an improvement and continue until it meets our specific criterion for matching. [King and Campbell used MINIPACK-1. (See articles by Moré et al.)]

In their demonstrations of this approach to analyze a [001] tilt GB in Nb, King and Campbell varied four parameters: thickness, defocus, x -tilt, and y -tilt. The steps were as follows

- They first optimized the electron-optical parameters using a 64×64 pixel image, giving $N = 4096$ and an image computational cell of 3.303 nm by 3.303 nm. Using the EMS program (Section 1.6), the optimization took 20 iterations and 80 multislice calculations.
- Next they had to optimize the structure of the GB. This process required defining 84 atomic positions in a unit cell of 4.16 nm \times 1.04 nm and a 512×128 (= 65,536) pixel image. Now the optimization required 16 iterations and 1300 multislice calculations.

These numbers are instructive. First they tell you that this computation can be done, which wasn't obvious. Second they tell you that this is a computer-intensive process; that we could have guessed!

You'll need to take enormous care in this type of analysis

- Align the simulated cell with the experimental cell and measure the unit cell in pixels.
- Choose a number of cells and relate them by the translation vector parallel to the rows of the image array.
- Calculate the standard-deviation images.
- Rotate the unit cell and repeat the exercise several times.

The orientation which gives the smallest standard deviation is your alignment. You must now adjust the magnification of the experimental image to fit the simulation, in a similar way to what you did for rotation. Next, you have to match the origins of both cells; the procedure is the same as we just described but translating the unit cell not rotating it. For a bicrystal, you now repeat this exercise for the other grain and then for the GB. You can improve the fit further if you take account of a constant background contribution, which probably arises due to the amorphous layer on both surfaces. Comparing experimental and calculated images quantitatively, we define f_i^{obs} as the corresponding value in the simulated image. We then calculate the residual $f_i(x)$ as follows

$$f_i(x) = \frac{(f_i^{\text{obs}} - (f_i^{\text{calc}}(x) + b^{\text{fit}}))}{W_i} \quad (31.15)$$

where b^{fit} is included as a free parameter in the optimization procedure. King and Campbell's calculations showed that W_i could be expressed as

$$W_i = \sigma_i^{\text{obs}} + 0.05f_i^{\text{obs}} \quad (31.16)$$

where σ_i^{obs} is the standard deviation of the i th pixel. Examples of the experimental, best fit, and normalized residuals are shown in Figure 31.24 for images from a $\Sigma = 5$, (310), [001] GB in Nb.

31.14 A FOURIER TECHNIQUE FOR QUANTITATIVE ANALYSIS

Möbus et al. proposed using what is referred to as an adaptive Fourier-filtering technique. The HRTEM image is digitized in the usual manner and then a special spatial-frequency filter is applied. This type of mask is designed for analyzing regions which contain defects.

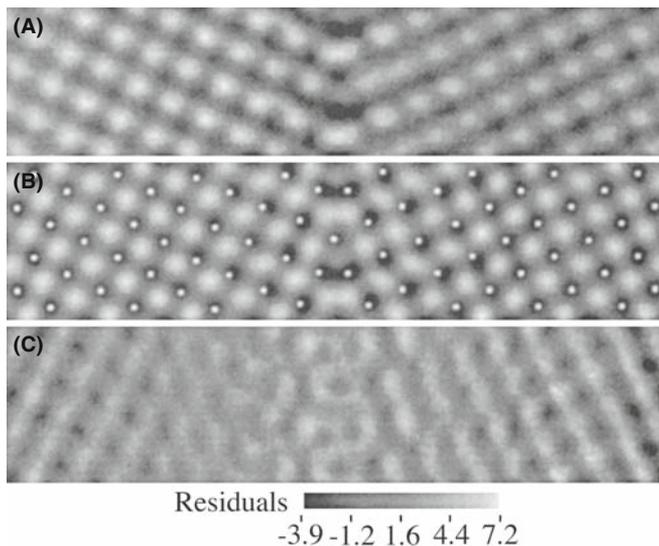


FIGURE 31.24. (A) Experimental image, (B) best fit simulation, and (C) normalized residuals of a $\Sigma = 5$ symmetric tilt boundary in Nb.

So the idea is that the computer automatically optimizes the mask to maximize the separation of the signal and the noise. This approach has not been widely practiced in TEM but clearly holds enormous promise. By varying the mask, this approach can prevent the analysis of a defect layer being dominated by the bulk information. Since the approach is quite straightforward signal processing, we will just illustrate an example found in the analysis of a simulated $\Sigma = 5$ GB with an extra period along the boundary. To test the analysis, white noise was added to a calculated image to give the image shown in Figure 31.25A. The power spectrum (the computer-generated DP) of the micrograph is shown in Figure 31.25B. The adaptive filter and the filtered

image are shown in Figure 31.25C and D. The important feature of the adaptive filter in this example is that it was created as such because the computer detected the doubling of the periodicity which is *only* present in the GB. Second the mask consists of elongated openings which we know we need when analyzing the GB because of the shape effect that we discussed when analyzing DPs (see Chapter 17).

AN ADAPTIVE FILTER
 One where the shape of the filter, or mask, is adapted to fit the shape of the 'image' it's filtering.

31.15 REAL OR RECIPROCAL SPACE?

In principle, we could equally well compare two images in reciprocal space rather than real space. However, because the Fourier transforms can generally be carried out much faster, the real-space approach has several advantages

- Fourier analysis separates local information into sine and cosine functions which are delocalized. When we reassemble the real-space image, higher parts of the frequency spectrum will be lost which will degrade the resolution.
- We want to maintain information on the absolute value of the intensities.
- Real-space methods are visually more intuitive for most of us. We can easily see what we have removed in the process.
- The real-space approach allows us to choose any values of n and m in defining our templates. Fourier space prefers aspect ratios given by 2^n .

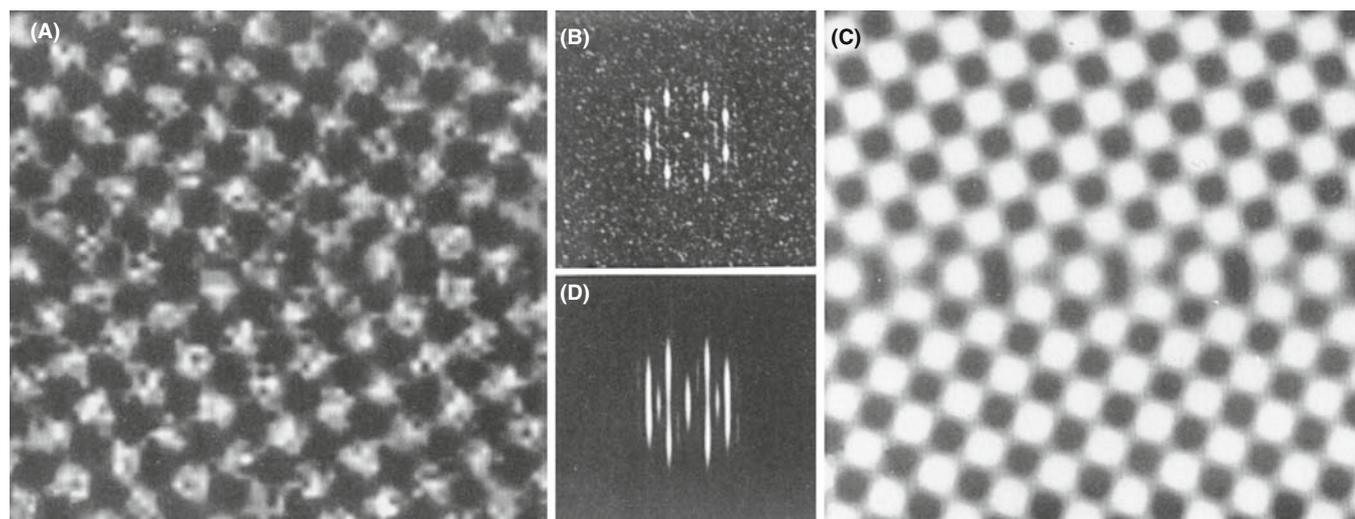


FIGURE 31.25. (A) White noise added to a calculated image of a $\Sigma = 5$ GB; (B) the power spectrum of (A); (C) the adaptive filter; (D) the filtered image.

31.16 SOFTWARE

Several software packages are widely used by the TEM community and have been listed in Section 1.6. We summarize the ones we know in Table 31.1. In its earliest application, image processing in TEM was almost exclusively applied to HRTEM images. This is no longer the case. Remember: always start with the best possible data. You can't always obtain a perfect image because your specimen might be beam sensitive, or coated with oxide, and you need to be aware of these limitations when processing or quantifying the image. This chapter has given you a hint of what is possible and where the subject is developing. We recommend that you obtain the software and start experimenting.

In this chapter we have discussed several different techniques used for processing TEM images. Different software is suitable for processing different types of image. The problem is that there may not be sufficient financial benefit unless a package supports another product; a similar problem was found for simulation software in Chapter 30. Many programs have been written

to run on particular computer systems and support can disappear even though the TEM has a lifetime of 15+ years.

31.17 THE OPTICAL BENCH—A LITTLE HISTORY

Although not widely used now, the optical bench is a useful instructional tool but you need to be old fashioned and record the image on film! A typical experimental setup is shown in Figure 31.26. The laser provides a coherent source of illumination representing the electron beam. The negative acts as the specimen. If it contains a set of lattice fringes, these act as a diffraction grating and give rise to a row of spots on the screen placed at the back focal plane of the 'objective' lens. The lens is thus performing an optical Fourier transform of the photograph. If you 'move the screen' to the image plane, the fringes reappear. You can make different masks and place them at the back focal plane or even create an 'adaptive filter' by exposing a photographic film and using this as the template for your mask. These masks correspond to the objective aperture in TEM. Students will find it instructive to transform their instructor or another suitable photograph, examine the frequency spectrum, and investigate the resulting spatial effect of different masks. The detail in the image is quickly lost as you remove the high spatial frequencies. This corresponds to inserting a smaller aperture in the back focal plane of the objective lens, as illustrated in Figure 31.27. So Figure 31.27D is effectively a BF image: clearly, you discard a lot of information in such images!

TABLE 31.1. Software for Image Processing

Package	Supplier	Cost
Digital Micrograph	Gatan	Commercial
NIH Image	NIH	Freeware

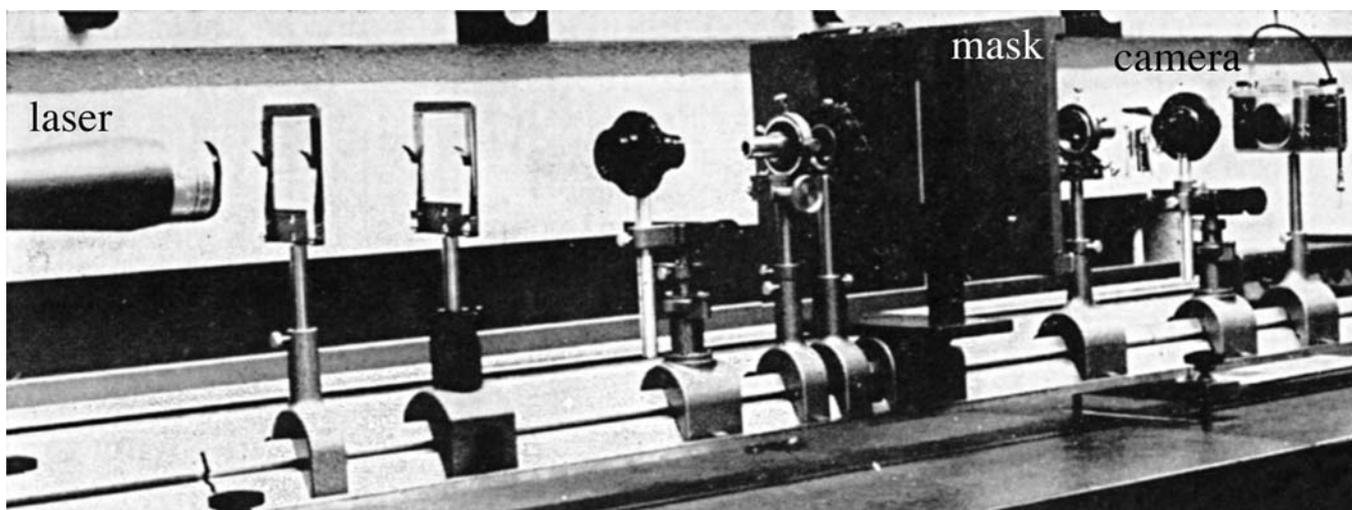


FIGURE 31.26. A typical experimental setup for an optical bench with the mask in the back focal plane.

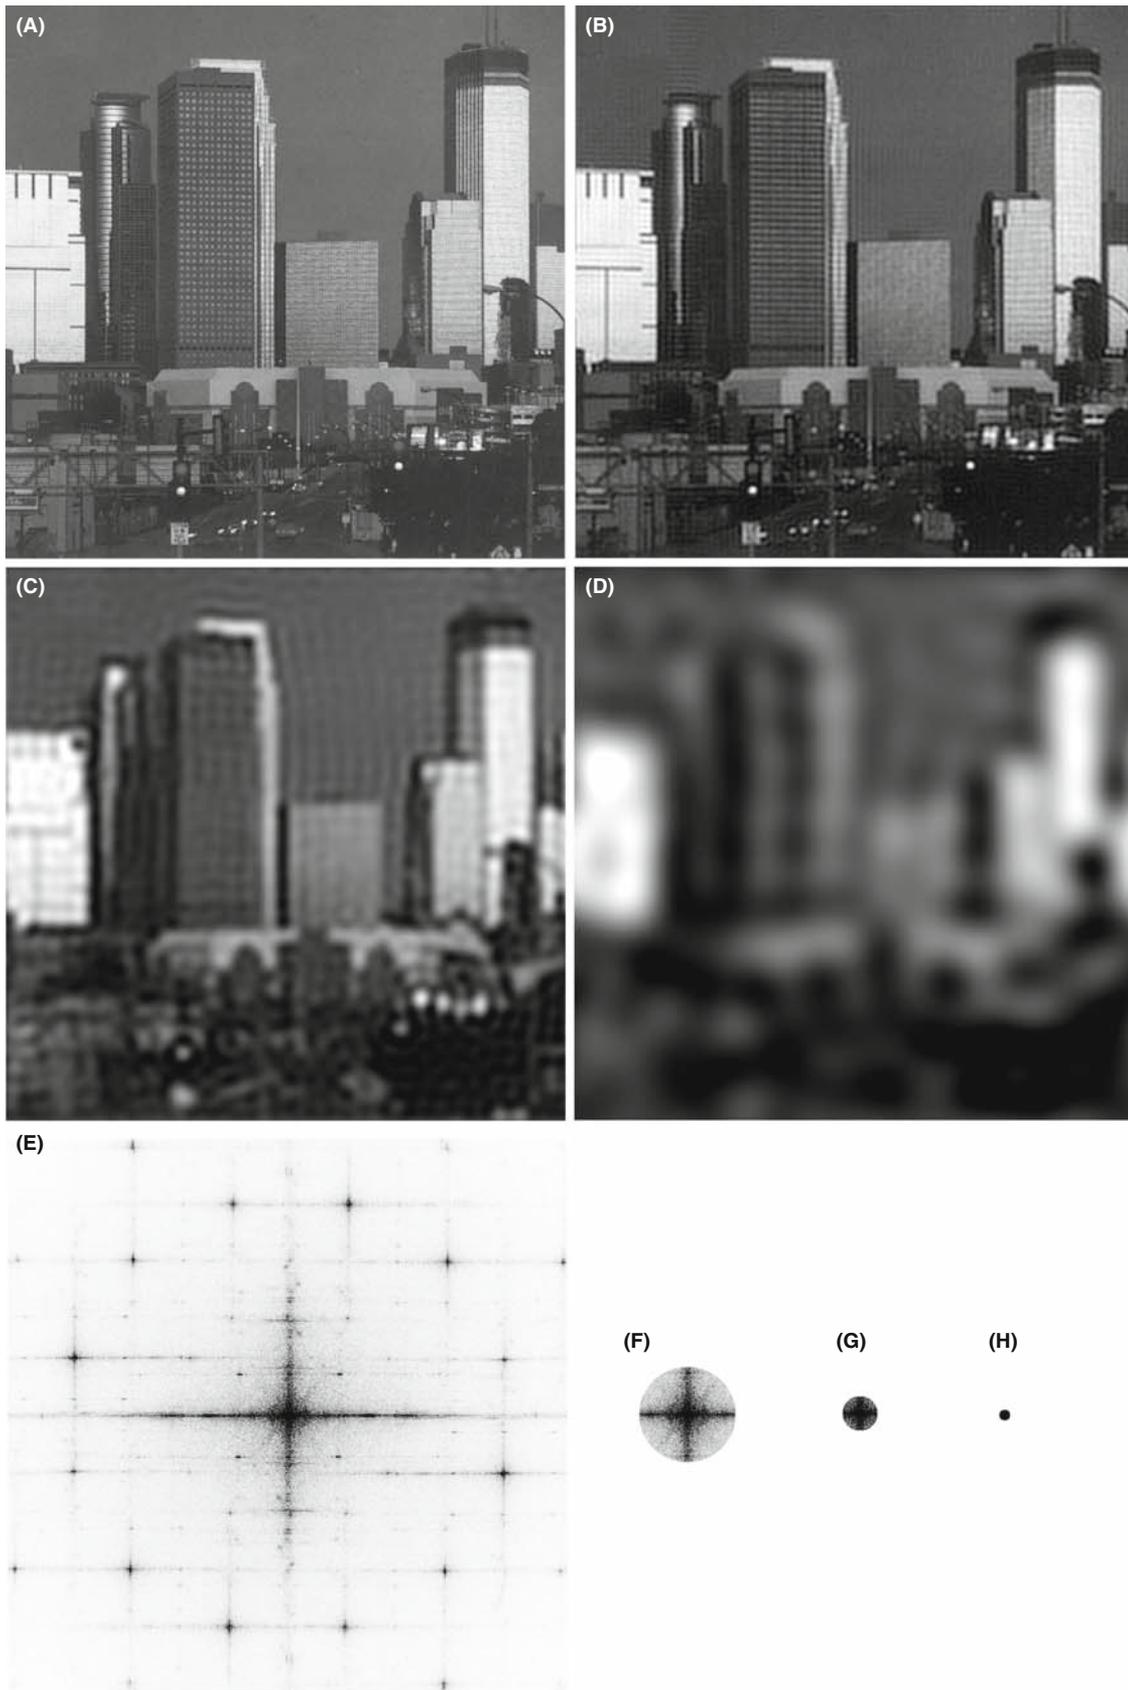


FIGURE 31.27. The effect of a mask (aperture) size on a non-periodic image of the Minneapolis skyline. (A–D) Reducing the aperture size, as indicated in the corresponding optical transform (diffractograms) (E–H) reduces the image detail. The streaks in (E) arise from the edges of the photographs.

CHAPTER SUMMARY

We have been doing image processing for many years; it's called 'dodging' in the photographic darkroom. You can even do this automatically with a special commercial enlarger. However, we have done very little quantitative imaging in materials science and even less in nanotechnology. The points you should remember when starting in the field are the following

- Quantitative comparison of simulated and experimental images depends on both the simulation program and the experimental parameters.
- If you are going to use reciprocal-space techniques for quantitative analysis, you should consider the shape of the optimum mask; it may not be a circular mask especially if you are studying interfaces! The process can, in principle, be assisted by the computer.
- The potential for image restoration is no longer limited by how the signals are mixed due to C_s and Δf . You can unscramble those effects if you have the right equipment.
- You will notice the repeated use of the word 'potential' where we don't mean $V(\mathbf{r})$! In many ways this chapter is a guide to the future of HRTEM and TEM in general. Some of the features won't be commonly available or optimized until the manufacturers realize their importance to the user and the user realizes that stability (for example) costs money.

Ultimately, the resolution limit is set by the signal-to-noise ratio in your image.

When you are processing images, there is always the possibility of removing information which is important. For example, Fresnel fringes often *should* be there! Beware of making reality match simulation, rather than the reverse. In the same vein, we draw your attention to the conclusion of Hÿtch and Stobbs, who found that they could only match their experimental and simulated images if they used a value for the specimen thickness which they knew was wrong! Their study emphasizes that, wherever possible, you should obtain independent measurements of the characteristics of your specimen and your machine. Remember the double-headed rhino in Figure 1.7; don't publish artifacts, even well-processed ones. Even if it's just for entertainment, images can reappear in unexpected places. Remember that an early image of Si dumbbells was later interpreted correctly because the authors gave the necessary data on how the image was formed when they published it. You should *always* do likewise.

PROCESSING IMAGES

Harrel, B *et al.* 1995 *Using Photoshop for Macintosh* Que Indianapolis. A detailed description of the possibilities using Photoshop; this volume should also serve as a warning for all microscopists. Ask the question: has the image you are examining been processed, and if so, how?

Hawkes, PW, Ed. 1980 *Computer Processing of Electron Microscope Images*, Springer-Verlag New York. Despite its age, this collection of articles is for advanced students.

ANALYZING C_s , Δf , ETC.

Krivanek, OL 1976 *Method for Determining the Coefficient of Spherical Aberration from a Single Electron Micrograph* *Optik* **45**, 97–101. An early paper giving a method for analyzing C_s and Δf using a diffractogram.

Zemlin, F, Weiss, K, Schiske, P, Kunath, W and Herrmann, K-H 1978 *Come-Free Alignment of High Resolution Electron Microscopes with the Aid of Optical Diffractograms* *Ultramicrosc.* **3** 49. The Zemlin tableau paper.

FOCAL SERIES AND TILT SERIES

Kirkland, EJ, and Siegel, BM, Uyeda, N and Fujiyoshi, Y 1982 *Nonlinear high resolution image processing of conventional transmission electron microscopy: II. Experiment* *Ultramicrosc.* **9**, 65–74.

Kirkland, AI, Saxton, WO, Chau, K-L, Tsuno, K and Kawasaki, M 1995 *Super-resolution by aperture synthesis: tilt series reconstruction in CTEM* *Ultramicrosc.* **57**, 355–374.

Kirkland, AI, Meyer, RR and Chang L-YS 2006 *Local Measurement and Computational Refinement of Aberrations for HR TEM* *Microsc. Microanal.* **12** 461–468.

Kirkland, AI and Meyer, RR 2004. *'Indirect' High-Resolution Transmission Electron Microscopy: Aberration Measurement and Wavefunction Reconstruction* Microsc. Microanal. **10** 401–413.

ADVANCED TEM PROCESSING

- Dorset, DL 1995 *Structural Electron Crystallography* Plenum New York. An early example of 3D reconstruction.
- Downing, KH 1992 Scanning Microscopy Supplement **6** AMF O'Hare IL p. 405.
- Hawkes, PW, Ed. 1992 *Signal and Image Processing in Microscopy and Microanalysis* Scanning Microscopy Supplement **6**, AMF O'Hare IL. This volume and its 1988 and 1996 companions are required reading for this subject.
- Hawkes, PW, Ottensmeyer, FP, Rosenfeld, A and Saxton, WO (Eds) 1988 *Image and Signal Processing for Electron Microscopy* Scanning Microscopy Supplement **2** AMF O'Hare IL.
- Hawkes, PW, Saxton, WO and Frank, J Eds. 1996 *Image Processing* Scanning Microscopy Supplement **X** AMF O'Hare IL.
- Russ, J.C. (1990) *Computer-Assisted Microscopy* Plenum Press New York. Chapter 3 is particularly relevant. Subsequent chapters give detailed analyses of the topics introduced here.
- Russ, JC 1995 *The Image Processing Handbook* 2nd edition CRC Press Boca Raton. A beautiful, comprehensive, and an essential component of any EM lab (or home).
- Saxton, WO 1992 Scanning Microscopy Supplement **6**, AMF O'Hare IL p. 405. Correcting for distortion in the specimen (Section 31.6B).
- Trus, BL, Unser, M, Pun, T and Stevens, AC 1992 Scanning Microscopy Supplement **6**, AMF O'Hare IL p. 441.

CHEMISTRY FROM IMAGES (SECTION 31.10)

- Kisielowski, C, Schwander, P, Baumann, FH, Seibt, M, Kim, Y, and Ourmazd, A 1995 *An Approach to Quantitative High-Resolution Transmission Electron Microscopy of Crystalline Materials* Ultramicrosc. **58**, 131–155.
- Ourmazd, A Baumann, FH, Bode, M and Kim, 1990 *Quantitative Chemical Lattice Imaging: Theory and Practice* Ultramicrosc. **34**, 237–255.

PATTERN RECOGNITION AND FILTERING ON INTERFACES

- King, WE and Campbell, GH 1993 *Determination of Thickness and Defocus by Quantitative Comparison of Experimental and Simulated High-Resolution Images* Ultramicrosc. **51**, 128–135.
- King, WE and Campbell, GH 1994 *Quantitative HREM Using Non-Linear Least-Squares Methods* Ultramicrosc. **56**, 46–53.
- Möbus, G, Necker, G and Rühle, M 1993 *Adaptive Fourier-Filtering Technique for Quantitative Evaluation of High-Resolution Electron Micrographs of Interfaces* Ultramicrosc. **49**, 46–65. The adaptive Fourier-filtering technique.
- Paciornik, S, Kilaas, R, Turner, J, and Dahmen, U 1995 *A Pattern Recognition Technique for the Analysis of Grain Boundary Structure by HREM* Ultramicrosc. **62**, 15–27. Pattern recognition at interfaces.

SPECIFICS

- Erickson, HP and Klug, A 1971 *Measurement and Compensation of Defocusing and Aberrations by Fourier Processing of Electron Micrographs* Phil. Trans. Roy. Soc. London B. **261**, 105–118. An example from the Nobel Prize winner.
- Frank, J 1980 in *Computer Processing of Electron Microscope Images*, Springer-Verlag New York. p. 187. Correlation techniques (Section 31.12).
- Hýtch, MJ and Stobbs, WM 1994 *Quantitative Comparison of High Resolution TEM images with image simulations* Ultramicrosc. **53** 191–203. A very interesting conclusion on the thickness of the specimen.
- Koster, AJ, van den Bos, A and van der Mast, KD 1988 Scanning Microscopy Supplement **2** AMF O'Hare IL p. 83. Autoalignment of the TEM.
- Koster, AJ and de Juijter, WJ 1992 *Practical Autoalignment of Transmission Electron Microscopes* Ultramicrosc. **40** 89–107. Autoalignment of the TEM.
- Krivanek, OL and Mooney, PE 1993 *Applications of Slow-Scan CCD Cameras in Transmission Electron Microscopy* Ultramicrosc. **49** 95–108. Autoalignment of the TEM.
- Moré, JJ 1977, in *Lectures Notes in Mathematics*, Ed. G.A. Watson p. 630 Springer Berlin.
- Moré, JJ, Garbow, BS and Hillstrom, KE 1980 *User Guide for MINIPACK-1*.
- Saxton, WO and Koch, TL 1982 *Interactive Image Processing with an Off-Line Minicomputer: Organization, Performance and Applications* J. Microsc. **127** 69–83. Autoalignment of the TEM.
- Saxton, WO, Smith, DJ and Erasmus, SJ 1983 *Procedures for Focusing, Stigmating and Alignment in High Resolution Electron Microscopy* J. Microsc. **130** 187–201. Cross-correlating pairs of images.

- Tang, D, Jansen, J, Zandbergen, HW and Schenk, H 1995 *The Estimation of Crystal Thickness and the Restoration of Structure-Factor Modulus from Electron Diffraction: A Kinematical Approach* Acta Cryst. **A51** 188–197. (31.6E) Quantitative analysis of DPs.
- van Heel, M, Winkler, H, Orlora, E and Schatz, M 1992 Scanning Microscopy Supplement **6**, AMF O'Hare IL p. 23. Review of TEM of beam-sensitive materials.
- Zou, XD 1995 *Electron Crystallography of Inorganic Structures*, Chemical Communications Stockholm University, Stockholm, Sweden.

THE COMPANION TEXT

The topic of image processing is often associated only with HRTEM images. This aspect is the subject of an entire chapter on focal-series reconstruction in the companion text.

One point you must always keep in mind is that the data are what you recorded originally. You will usually modify these data in some way before it is reported to the scientific community. The question you must always keep in mind is how much detail on the processing must you report so that other scientists will be able to check and repeat your analysis.

SELF-ASSESSMENT QUESTIONS

- Q31.1 Define the term 'quantifying HRTEM images.'
- Q31.2 Why is Fourier filtering so named?
- Q31.3 Sketch a pair of diffractograms from an amorphous Ge film with $\Delta f = 1$ sch and 4 sch, respectively.
- Q31.4 Why is amorphous Ge used rather than amorphous C to support Au nanoparticles in testing the resolution of an HRTEM?
- Q31.5 If we plot $\frac{n}{u^2}$ against u^2 , what is the slope of the curve and what is n ?
- Q31.6 What is unsharp masking?
- Q31.7 How would you obtain the best TEM still image from a TEM video?
- Q31.8 What is a kernel in image processing?
- Q31.9 Give an example of a kernel for enhancing the contrast in image processing.
- Q31.10 Why do we say that the signal-to-noise ratio will always be small in low-dose TEM?
- Q31.11 Why is it helpful to record a focal series of images if you want to do quantitative HRTEM?
- Q31.12 Will Kikuchi diffraction influence electron crystallography using DPs?
- Q31.13 What is the reason for using a tilted-beam series to enhance resolution in HRTEM?
- Q31.14 What is the most important use of remote microscopy?
- Q31.15 We use a diffractogram tableau to bring the beam alignment to within 0.1 mrad of the optic axis. What aberration are we correcting?
- Q31.16 What is the use of pattern recognition in HRTEM and why would you like to have a new 4 k × 4 k camera for this purpose?
- Q31.17 What is the basic principle behind QUANTITEM?
- Q31.18 There are two methods for estimating the 'goodness of fit'? What are they and how do they differ?
- Q31.19 For the best quantitative data, you should use a focal series and low-dose imaging. What fundamental principle of quantum physics does this statement remind you of?
- Q31.20 What is the adaptive Fourier-filtering technique?

TEXT-SPECIFIC QUESTIONS

Some of the questions for this chapter require that you have access to image-processing software.

- T31.1 Identify another particle/matrix pair that might be studied like that in Figures 31.2 and 31.3. Describe how you would use image processing in your analysis to show it is a spherical particle, not a cylinder.
- T31.2 Describe how you would extract the best image from a 30-second video clip of a stable grain boundary. (Hint: use several sections in this chapter.)
- T31.3 Use your scanner and Digital Micrograph to check the correlation between the images and diffractograms in Figures 31.4 and 31.5. You can consider the original publications. Comment on your results and the scientific process.
- T31.4 By considering the original paper or just from scratch, reproduce the graphs shown in Figure 31.6A.
- T31.5 Explain how you would actually use the kernel for a digital Laplacian in analyzing a TEM image. What type of image might this be useful for? Find (in the literature) and describe two other kernels.
- T31.6 Low-dose microscopy should be used much more routinely. Describe in detail the routines that are actually used by two different cryo-TEM manufacturers.
- T31.7 Consider Figure 30.10. Summarize how you would carry out this type of study in your lab. Pay special attention to the important features of the specimen (e.g., is it applicable to all specimens).

- T31.8 By delving into the literature, write a description of how you would use a program like QUANTITEM. Identify other groups who have used similar programs. Give a critical review of the applications of the concept of parameterization in TEM (with special attention to the range of materials that can be studied).
- T31.9 Use a laser pointer and a small-mesh TEM grid to produce a diffraction pattern of the grid on a nearby wall. (See Figure 31.27.) Relate the spacing of the spots to the distances (e.g., mesh spacing) that you used.
- T31.10 Use Digital Micrograph, or a similar product, to reproduce the images shown in Figure 31.27 using your home town or a face as the object. Comment on the reduction of resolution as you decrease the aperture diameter.