7 DIGITAL IMAGE PROCESSING

7.1 INTRODUCTION

Digital image processing involves the manipulation and interpretation of digital images (Section 1.5) with the aid of a computer. This form of remote sensing actually began in the 1960s with a limited number of researchers analyzing airborne multispectral scanner data and digitized aerial photographs. However, it was not until the launch of Landsat-1, in 1972, that digital image data became widely available for land remote sensing applications. At that time, not only was the theory and practice of digital image processing in its infancy, but also the cost of digital computers was very high and their computational efficiency was very low by modern standards. Today, access to low cost, efficient computer hardware and software is commonplace, and the sources of digital image data are many and varied. These sources range from commercial and governmental earth resource satellite systems, to the meteorological satellites, to airborne scanner data, to airborne digital camera data, to image data generated by photogrammetric scanners and other high resolution digitizing systems. All of these forms of data can be processed and analyzed using the techniques described in this chapter.

Digital image processing is an extremely broad subject, and it often involves procedures that can be mathematically complex. Our objective in this

chapter is to introduce the basic principles of digital image processing without delving into the detailed mathematics involved. Also, we avoid extensive treatment of the rapidly changing hardware associated with digital image processing. The references at the end of this chapter are provided for those wishing to pursue such additional detail.

The central idea behind digital image processing is quite simple. The digital image is fed into a computer one pixel at a time. The computer is programmed to insert these data into an equation, or series of equations, and then store the results of the computation for each pixel. These results form a new digital image that may be displayed or recorded in pictorial format or may itself be further manipulated by additional programs. The possible forms of digital image manipulation are literally infinite. However, virtually all these procedures may be categorized into one (or more) of the following seven broad types of computer-assisted operations:

1. **Image rectification and restoration.** These operations aim to correct distorted or degraded image data to create a more faithful representation of the original scene. This typically involves the initial processing of raw image data to correct for geometric distortions, to calibrate the data radiometrically, and to eliminate noise present in the data. Thus, the nature of any particular image restoration process is highly dependent upon the characteristics of the sensor used to acquire the image data. Image rectification and restoration procedures are often termed *preprocessing* operations because they normally precede further manipulation and analysis of the image data to extract specific information. We briefly discuss these procedures in Section 7.2 with treatment of various geometric corrections, radiometric corrections, and noise corrections.
2. **Image enhancement.** These procedures are applied to image data in order to more effectively display or record the data for subsequent visual interpretation. Normally, image enhancement involves techniques for increasing the visual distinctions between features in a scene. The objective is to create "new" images from the original image data in order to increase the amount of information that can be visually interpreted from the data. The enhanced images can be displayed interactively on a monitor or they can be recorded in a hardcopy format, either in black and white or in color. There are no simple rules for producing the single "best" image for a particular application. Often several enhancements made from the same "raw" image are necessary. We summarize the various broad approaches to enhancement in Section 7.3. In Section 7.4, we treat specific procedures that manipulate the contrast of an image (level slicing and contrast stretching). In Section 7.5, we discuss spatial feature manipulation (spatial filtering, convolution, edge enhancement, and Fourier analysis). In Section 7.6, we consider en-

hancements involving multiple spectral bands of imagery (spectral ratioing, principal and canonical components, vegetation components, and intensity–hue–saturation color space transformations).

3. **Image classification.** The objective of these operations is to replace visual analysis of the image data with quantitative techniques for automating the identification of features in a scene. This normally involves the analysis of multispectral image data and the application of statistically based decision rules for determining the land cover identity of each pixel in an image. When these decision rules are based solely on the spectral radiances observed in the data, we refer to the classification process as *spectral pattern recognition*. In contrast, the decision rules may be based on the geometric shapes, sizes, and patterns present in the image data. These procedures fall into the domain of *spatial pattern recognition*. In either case, the intent of the classification process is to categorize all pixels in a digital image into one of several land cover classes, or "themes." These categorized data may then be used to produce *thematic maps* of the land cover present in an image and/or produce summary statistics on the areas covered by each land cover type. Due to their importance, image classification procedures comprise the subject of more than one-third of the material in this chapter (Sections 7.7 to 7.16). We emphasize *spectral* pattern recognition procedures because the current state-of-the-art for these procedures is more advanced than for *spatial* pattern recognition approaches. (Substantial research is ongoing in the development of spatial and combined spectral/spatial image classification.) We emphasize "supervised," "unsupervised," and "hybrid" approaches to spectrally based image classification. We also describe various procedures for assessing the accuracy of image classification results.
4. **Data merging and GIS integration.** These procedures are used to combine image data for a given geographic area with other geographically referenced data sets for the same area. These other data sets might simply consist of image data generated on other dates by the same sensor or by other remote sensing systems. Frequently, the intent of data merging is to combine remotely sensed data with other sources of information in the context of a GIS. For example, image data are often combined with soil, topographic, ownership, zoning, and assessment information. We discuss data merging in Section 7.17. In this section, we highlight multitemporal data merging, change detection procedures, and multisensor image merging. We also illustrate the incorporation of GIS data in land cover classification.
5. **Hyperspectral image analysis.** Virtually all of the image processing principles introduced in this chapter in the context of multispectral image analysis may be extended directly to the analysis of

hyperspectral data. However, the basic nature and sheer volume of hyperspectral data sets is such that various image processing procedures have been developed to analyze such data specifically. We introduce these procedures in Section 7.18.

6. **Biophysical modeling.** The objective of biophysical modeling is to relate quantitatively the digital data recorded by a remote sensing system to biophysical features and phenomena measured on the ground. For example, remotely sensed data might be used to estimate such varied parameters as crop yield, pollution concentration, or water depth. Likewise, remotely sensed data are often used in concert with GIS techniques to facilitate environmental modeling. The intent of these operations is to simulate the functioning of environmental systems in a spatially explicit manner and to predict their behavior under altered ("what if") conditions, such as global climate change. We discuss the overall concept of biophysical modeling in Section 7.19. In Section 7.20 we highlight the important role spatial resolution plays in the incorporation of remotely sensed data in the environmental modeling process.
7. **Image transmission and compression.** Given the increasingly high volume of data available from remote sensing systems and the distribution of image data over the Internet, image compression techniques are the subject of continuing image processing research. We briefly introduce the use of wavelet transforms in this context in Section 7.21.

We have made the above subdivisions of the topic of digital image processing to provide the reader with a conceptual roadmap for studying this chapter. Although we treat each of these procedures as distinct operations, they all interrelate. For example, the restoration process of noise removal can often be considered an enhancement procedure. Likewise, certain enhancement procedures (such as principal components analysis) can be used not only to enhance the data but also to improve the efficiency of classification operations. In a similar vein, data merging can be used in image classification in order to improve classification accuracy. Hence, the boundaries between the various operations we discuss separately here are not well defined in practice.

7.2 IMAGE RECTIFICATION AND RESTORATION

As previously mentioned, the intent of image rectification and restoration is to correct image data for distortions or degradations that stem from the image acquisition process. Obviously, the nature of such procedures varies considerably with such factors as the digital image acquisition type (digital camera, along-track scanner, across-track scanner), platform (airborne versus satellite), and total field of view. We make no attempt to describe the entire

range of image rectification and restoration procedures applied to each of these various types of factors. Rather, we treat these operations under the generic headings of geometric correction, radiometric correction, and noise removal.

Geometric Correction

Raw digital images usually contain geometric distortions so significant that they cannot be used directly as a map base without subsequent processing. The sources of these distortions range from variations in the altitude, attitude, and velocity of the sensor platform to factors such as panoramic distortion, earth curvature, atmospheric refraction, relief displacement, and nonlinearities in the sweep of a sensor's IFOV. The intent of geometric correction is to compensate for the distortions introduced by these factors so that the corrected image will have the highest practical geometric integrity.

The geometric correction process is normally implemented as a two-step procedure. First, those distortions that are *systematic*, or predictable, are considered. Second, those distortions that are essentially *random*, or unpredictable, are considered.

Systematic distortions are well understood and easily corrected by applying formulas derived by modeling the sources of the distortions mathematically. For example, a highly systematic source of distortion involved in multispectral scanning from satellite altitudes is the eastward rotation of the earth beneath the satellite during imaging. This causes each optical sweep of the scanner to cover an area slightly to the west of the previous sweep. This is known as *skew distortion*. The process of *deskewing* the resulting imagery involves offsetting each successive scan line slightly to the west. The skewed-parallelogram appearance of satellite multispectral scanner data is a result of this correction.

Random distortions and residual unknown systematic distortions are corrected by analyzing well-distributed ground control points (GCPs) occurring in an image. As with their counterparts on aerial photographs, GCPs are features of known ground location that can be accurately located on the digital imagery. Some features that might make good control points are highway intersections and distinct shoreline features. In the correction process numerous GCPs are located both in terms of their two image coordinates (column, row numbers) on the distorted image and in terms of their ground coordinates (typically measured from a map, or GPS located in the field, in terms of UTM coordinates or latitude and longitude). These values are then submitted to a least squares regression analysis to determine coefficients for two *coordinate transformation equations* that can be used to interrelate the geometrically correct (map) coordinates and the distorted-image coordinates. (Appendix C describes one of the more common forms of coordinate transformation, the

affine transformation.) Once the coefficients for these equations are determined, the distorted-image coordinates for any map position can be precisely estimated. Expressing this in mathematic notation,

$$x = f_1(X, Y) \qquad y = f_2(X, Y) \tag{7.1}$$

where

(x, y)= distorted-image coordinates (column, row)
(X, Y) = correct (map) coordinates
f_1, f_2 = transformation functions

Intuitively, it might seem as though the above equations are stated backward! That is, they specify how to determine the distorted-image positions corresponding to correct, or undistorted, map positions. But that is exactly what is done during the geometric correction process. We first define an undistorted output matrix of "empty" map cells and then fill in each cell with the gray level of the corresponding pixel, or pixels, in the distorted image. This process is illustrated in Figure 7.1. This diagram shows the geometrically

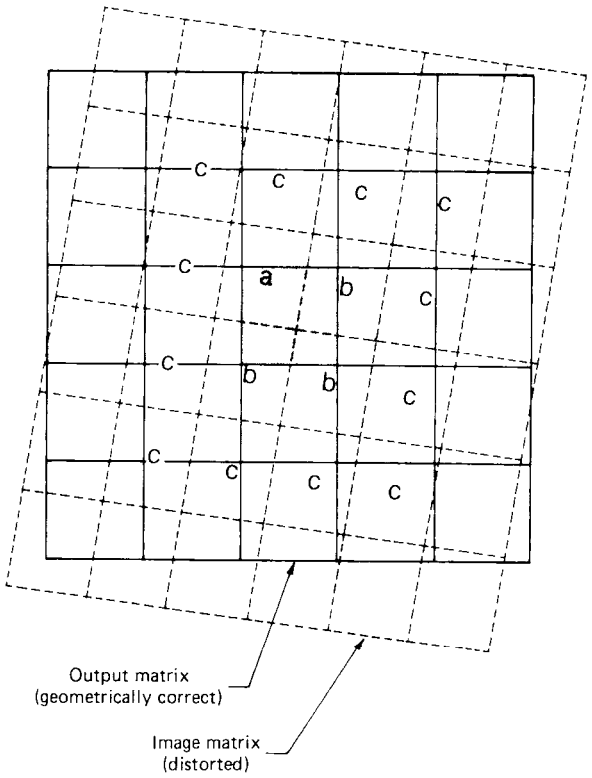

Figure 7.1 Matrix of geometrically correct output pixels superimposed on matrix of original, distorted input pixels.

correct output matrix of cells (solid lines) superimposed over the original, distorted matrix of image pixels (dashed lines). After producing the transformation function, a process called *resampling* is used to determine the pixel values to fill into the output matrix from the original image matrix. This process is performed using the following operations:

1. The coordinates of each element in the undistorted output matrix are transformed to determine their corresponding location in the original input (distorted-image) matrix.
2. In general, a cell in the output matrix will not directly overlay a pixel in the input matrix. Accordingly, the intensity value or digital number (DN) eventually assigned to a cell in the output matrix is determined on the basis of the pixel values that surround its transformed position in the original input matrix.

A number of different resampling schemes can be used to assign the appropriate DN to an output cell or pixel. To illustrate this, consider the shaded output pixel shown in Figure 7.1. The DN for this pixel could be assigned simply on the basis of the DN of the closest pixel in the input matrix, disregarding the slight offset. In our example, the DN of the input pixel labeled *a* would be transferred to the shaded output pixel. This approach is called *nearest neighbor* resampling. It offers the advantage of computational simplicity and avoids having to alter the original input pixel values. However, features in the output matrix may be offset spatially by up to one-half pixel. This can cause a disjointed appearance in the output image product. Figure 7.2*b* is an example of a nearest neighbor resampled Landsat TM image. Figure 7.2*a* shows the original, distorted image.

More sophisticated methods of resampling evaluate the values of several pixels surrounding a given pixel in the input image to establish a "synthetic" DN to be assigned to its corresponding pixel in the output image. The *bilinear interpolation* technique takes a distance-weighted average of the DNs of the four nearest pixels (labeled *a* and *b* in the distorted-image matrix in Figure 7.1). This process is simply the two-dimensional equivalent to linear interpolation. As shown in Figure 7.2*c*, this technique generates a smoother appearing resampled image. However, because the process alters the gray levels of the original image, problems may be encountered in subsequent spectral pattern recognition analyses of the data. (Because of this, resampling is often performed after, rather than prior to, image classification procedures.)

An improved restoration of the image is provided by the *bicubic interpolation* or *cubic convolution* method of resampling. In this approach, the transferred synthetic pixel values are determined by evaluating the block of 16 pixels in the input matrix that surrounds each output pixel (labeled *a*, *b*, and *c* in Figure 7.1). Cubic convolution resampling (Figure 7.2*d*) avoids the disjointed appearance of the nearest neighbor method and provides a slightly

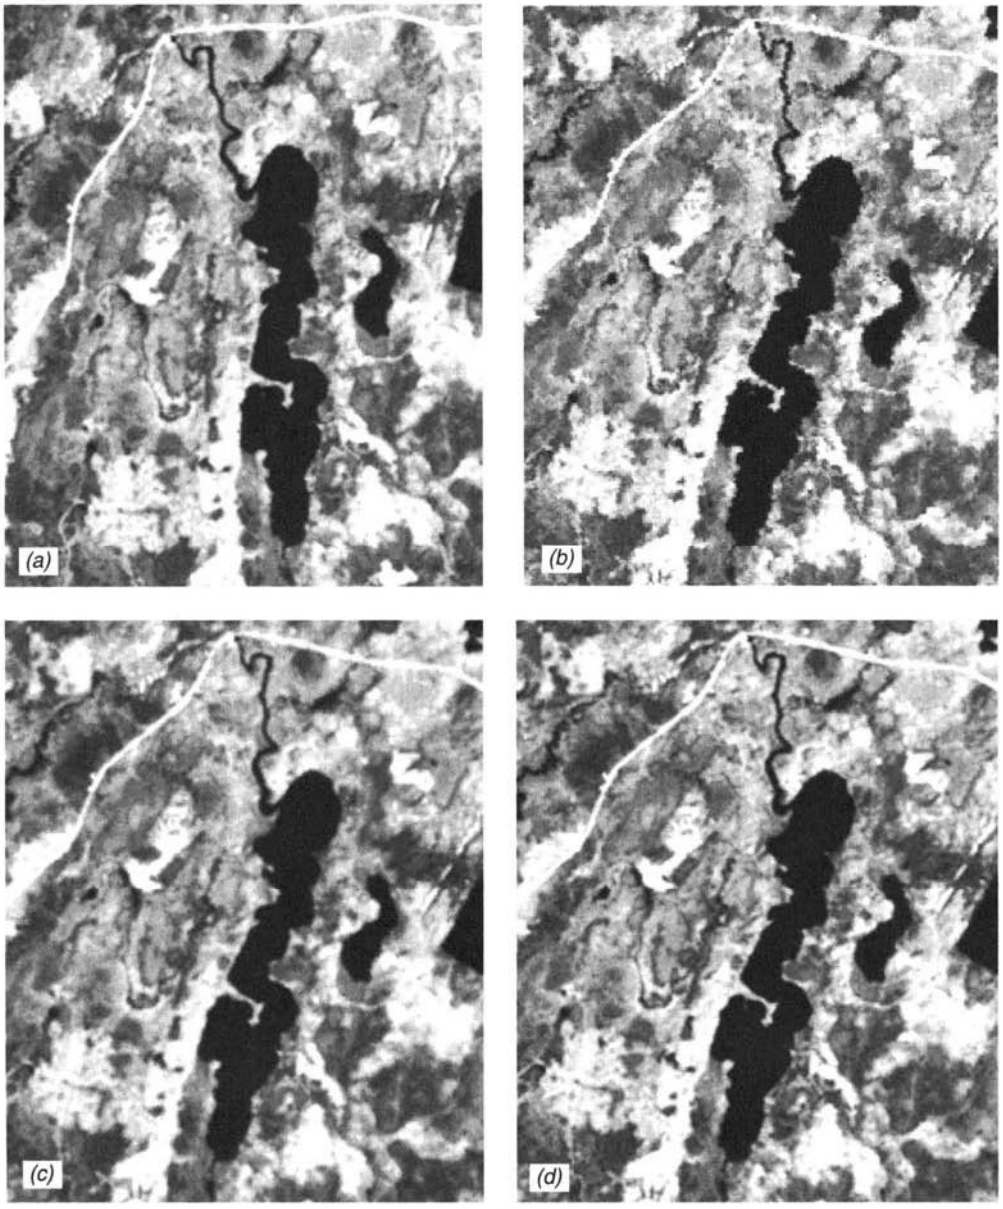

Figure 7.2 Resampling results: (*a*) original Landsat TM data; (*b*) nearest neighbor assignment; (*c*) bilinear interpolation; (*d*) cubic convolution. Scale 1:100,000.

sharper image than the bilinear interpolation method. (Again, this method alters the original image gray levels to some extent and other types of resampling can be employed to minimize this effect.)

As we discuss later, resampling techniques are important in several digital processing operations besides the geometric correction of raw images. For example, resampling is used to overlay or register multiple dates of imagery. It is also used to register images of differing resolution. Also, resampling procedures are used extensively to register image data and other sources of data in GISs. (Recall that resampling was discussed as an important part of digital orthophoto production in Section 3.9.) Appendix C contains additional details about the implementation of the various resampling procedures discussed in this section.

Radiometric Correction

As with geometric correction, the type of radiometric correction applied to any given digital image data set varies widely among sensors. Other things being equal, the radiance measured by any given system over a given object is influenced by such factors as changes in scene illumination, atmospheric conditions, viewing geometry, and instrument response characteristics. Some of these effects, such as viewing geometry variations, are greater in the case of airborne data collection than in satellite image acquisition. Also, the need to perform correction for any or all of these influences depends directly upon the particular application at hand.

In the case of satellite sensing in the visible and near-infrared portion of the spectrum, it is often desirable to generate mosaics of images taken at different times or to study the changes in the reflectance of ground features at different times or locations. In such applications, it is usually necessary to apply a *sun elevation correction* and an *earth–sun distance correction*. The sun elevation correction accounts for the seasonal position of the sun relative to the earth (Figure 7.3). Through this process, image data acquired under different solar illumination angles are normalized by calculating pixel brightness values assuming the sun was at the zenith on each date of sensing. The correction is usually applied by dividing each pixel value in a scene by the sine of the solar elevation angle for the particular time and location of imaging. Alternatively, the correction is applied in terms of the sun's angle from the zenith, which is simply 90° minus the solar elevation angle. (In this case each pixel value is divided by the cosine of the sun's angle from the zenith, resulting in the identical correction.) In either case, the correction ignores topographic and atmospheric effects.

The earth–sun distance correction is applied to normalize for the seasonal changes in the distance between the earth and the sun. The earth–sun distance

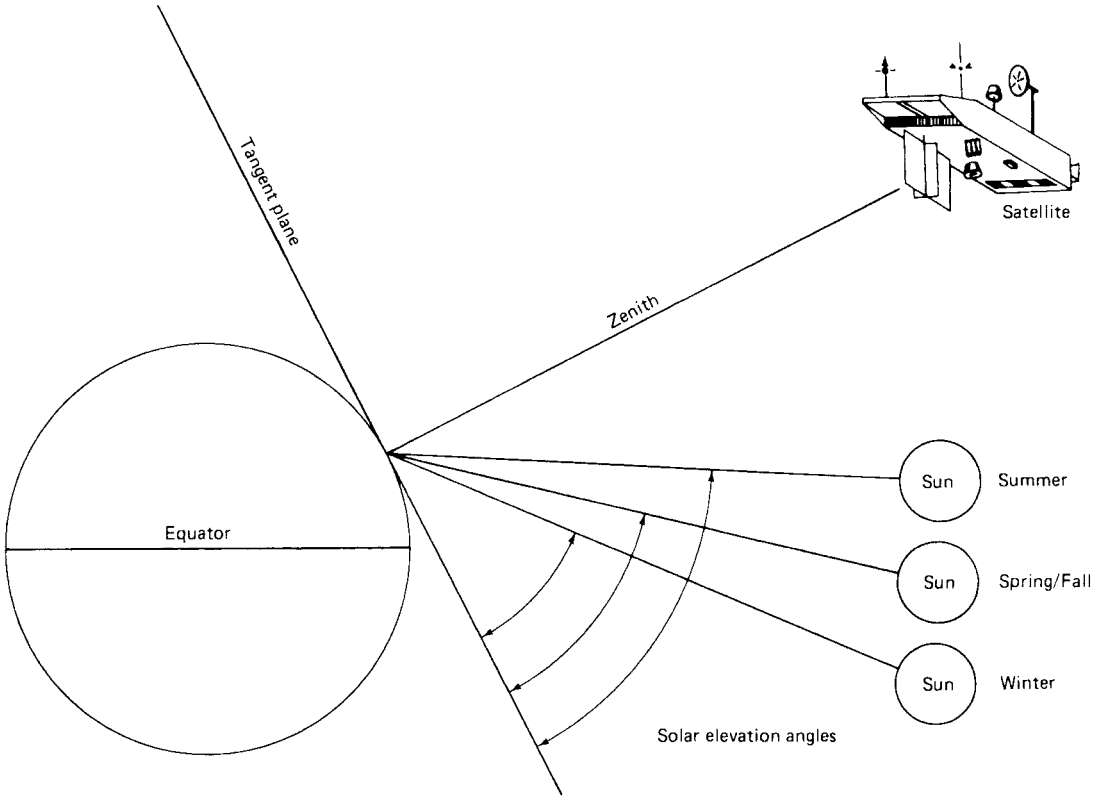

Figure 7.3 Effects of seasonal change on solar elevation angle. (The solar zenith angle is equal to 90° minus the solar elevation angle.)

is usually expressed in astronomical units. (An astronomical unit is equivalent to the mean distance between the earth and the sun, approximately 149.6 × 10^6 km.) The irradiance from the sun decreases as the square of the earth–sun distance.

Ignoring atmospheric effects, the combined influence of solar zenith angle and earth–sun distance on the irradiance incident on the earth's surface can be expressed as

$$E = \frac{E_0 \cos \theta_0}{d^2} \tag{7.2}$$

where

E = normalized solar irradiance
E_0 = solar irradiance at mean earth–sun distance
θ_0 = sun's angle from the zenith
d = earth–sun distance, in astronomical units

(Information on the solar elevation angle and earth–sun distance for a given scene is normally part of the ancillary data supplied with the digital data.)

As initially discussed in Section 1.4, the influence of solar illumination variation is compounded by atmospheric effects. The atmosphere affects the radiance measured at any point in the scene in two contradictory ways. First, it attenuates (reduces) the energy illuminating a ground object. Second, it acts as a reflector itself, adding a scattered, extraneous "path radiance" to the signal detected by a sensor. Thus, the composite signal observed at any given pixel location can be expressed by

$$L_{tot} = \frac{\rho E T}{\pi} + L_p \tag{7.3}$$

where

L_{tot} = total spectral radiance measured by sensor
ρ = reflectance of object
E = irradiance on object
T = transmission of atmosphere
L_p = path radiance

(All of the above quantities depend on wavelength.)

Only the first term in the above equation contains valid information about ground reflectance. The second term represents the scattered path radiance, which introduces "haze" in the imagery and reduces image contrast. (Recall that scattering is wavelength dependent, with shorter wavelengths normally manifesting greater scattering effects.) *Haze compensation* procedures are designed to minimize the influence of path radiance effects. One means of haze compensation in multispectral data is to observe the radiance recorded over target areas of essentially zero reflectance. For example, the reflectance of deep clear water is essentially zero in the near-infrared region of the spectrum. Therefore, any signal observed over such an area represents the path radiance, and this value can be subtracted from all pixels in that band.

For convenience, haze correction routines are often applied uniformly throughout a scene. This may or may not be valid, depending on the uniformity of the atmosphere over a scene. When extreme viewing angles are involved in image acquisition, it is often necessary to compensate for the influence of varying the atmospheric path length through which the scene is recording. In such cases, off-nadir pixel values are usually normalized to their nadir equivalents.

Another radiometric data processing activity involved in many quantitative applications of digital image data is *conversion of DNs to absolute radiance values*. This operation accounts for the exact form of the A-to-D response functions for a given sensor and is essential in applications where

measurement of absolute radiances is required. For example, such conversions are necessary when changes in the absolute reflectance of objects are to be measured over time using different sensors (e.g., the multispectral scanner on Landsat-3 versus that on Landsat-5). Likewise, such conversions are important in the development of mathematical models that physically relate image data to quantitative ground measurements (e.g., water quality data).

Normally, detectors and data systems are designed to produce a linear response to incident spectral radiance. For example, Figure 7.4 shows the linear radiometric response function typical of an individual TM channel. Each spectral band of the sensor has its own response function, and its characteristics are monitored using onboard calibration lamps (and temperature references for the thermal channel). The absolute spectral radiance output of the calibration sources is known from prelaunch calibration and is assumed to be stable over the life of the sensor. Thus, the onboard calibration sources form the basis for constructing the radiometric response function by relating known radiance values incident on the detectors to the resulting DNs.

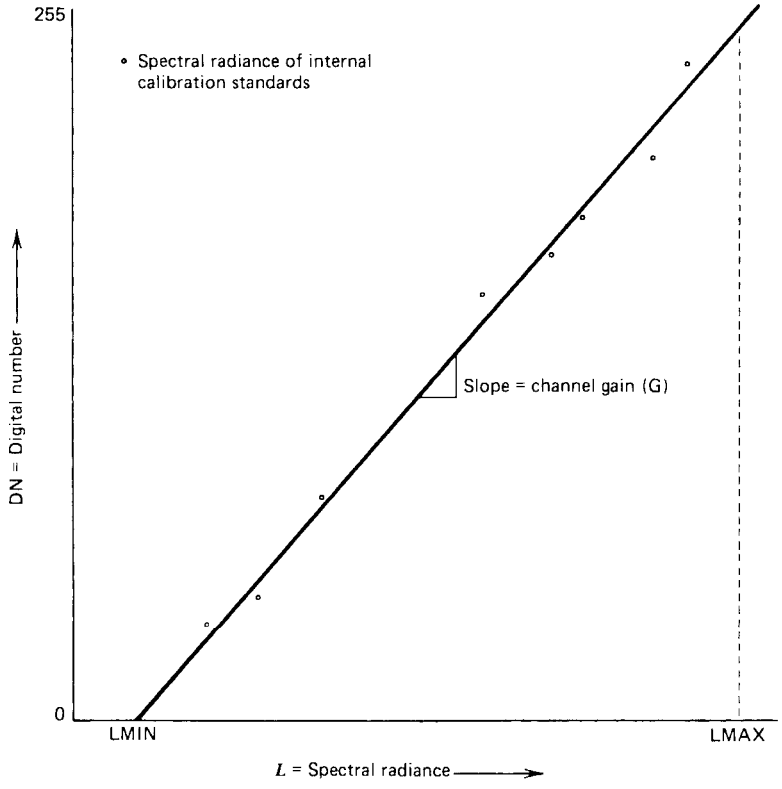

Figure 7.4 Radiometric response function for an individual TM channel.

It can be seen from Figure 7.4 that a linear fit to the calibration data results in the following relationship between radiance and DN values for any given channel:

$$\text{DN} = GL + B \tag{7.4}$$

where

DN = digital number value recorded
G = slope of response function (channel gain)
L = spectral radiance measured (over the spectral bandwidth of the channel)
B = intercept of response function (channel offset)

Note that the slope and intercept of the above function are referred to as the *gain* and *offset* of the response function, respectively. In Figure 7.4 LMIN is the spectral radiance corresponding to a DN response of 0 and LMAX is the minimum radiance required to generate the maximum DN (here 255). That is, LMAX represents the radiance at which the channel saturates. The range from LMIN to LMAX is the dynamic range for the channel.

Figure 7.5 is a plot of the inverse of the radiometric response. Here we have simply interchanged the axes from Figure 7.4. The equation for this line is

$$L = \left(\frac{\text{LMAX} - \text{LMIN}}{255}\right) \text{DN} + \text{LMIN} \tag{7.5}$$

Equation 7.5 can be used to convert any DN in a particular band to absolute units of spectral radiance in that band if LMAX and LMIN are known from the sensor calibration.

Often the LMAX and LMIN values published for a given sensor are expressed in units of mW cm^{-2} sr^{-1} μm^{-1}. That is, the values are often specified in terms of radiance per unit wavelength. To estimate the total within-band radiance in such cases, the value obtained from Eq. 7.5 must be multiplied by the width of the spectral band under consideration. Hence, a precise estimate of within-band radiance requires detailed knowledge of the spectral response curves for each band.

Noise Removal

Image noise is any unwanted disturbance in image data that is due to limitations in the sensing, signal digitization, or data recording process. The potential sources of noise range from periodic drift or malfunction of a detector, to electronic interference between sensor components, to intermittent "hiccups" in the data transmission and recording sequence. Noise can either degrade or totally mask the true radiometric information content of a digital image.

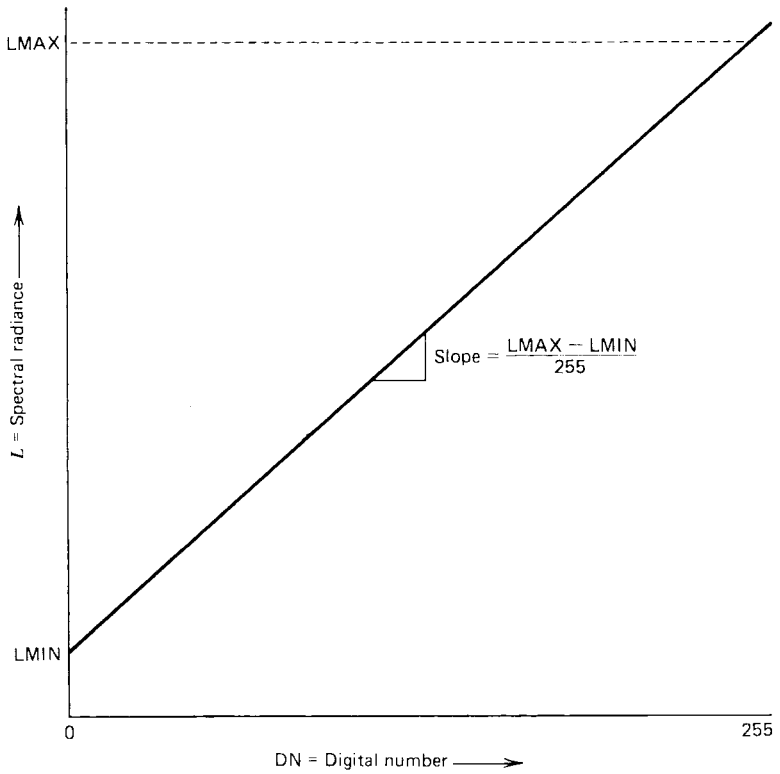

Figure 7.5 Inverse of radiometric response function for an individual TM channel.

Hence, noise removal usually precedes any subsequent enhancement or classification of the image data. The objective is to restore an image to as close an approximation of the original scene as possible.

As with geometric restoration procedures, the nature of noise correction required in any given situation depends upon whether the noise is systematic (periodic), random, or some combination of the two. For example, multispectral scanners that sweep multiple scan lines simultaneously often produce data containing systematic *striping* or *banding*. This stems from variations in the response of the individual detectors used within each band. Such problems were particularly prevalent in the collection of early Landsat MSS data. While the six detectors used for each band were carefully calibrated and matched prior to launch, the radiometric response of one or more tended to drift over time, resulting in relatively higher or lower values along every sixth line in the image data. In this case valid data are present in the defective lines, but they must be normalized with respect to their neighboring observations.

Several *destriping* procedures have been developed to deal with the type of problem described above. One method is to compile a set of histograms for

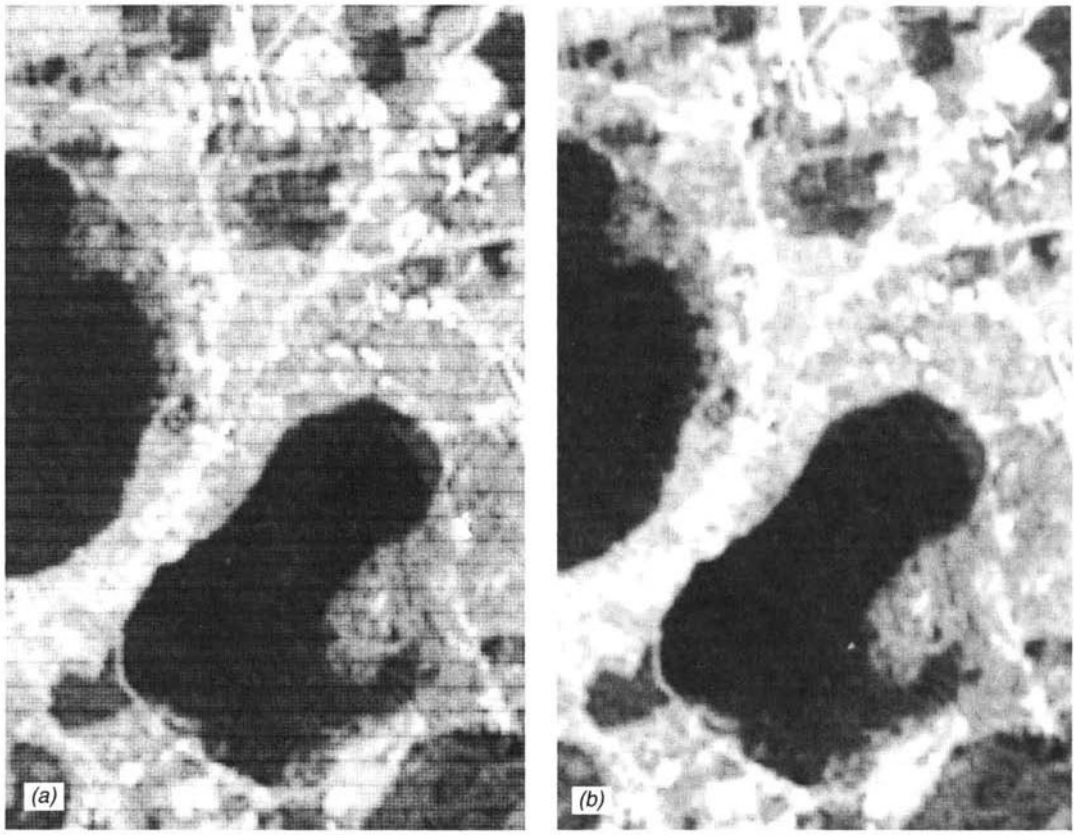

Figure 7.6 Destriping algorithm illustration: (*a*) original image manifesting striping with a six-line frequency; (*b*) restored image resulting from applying histogram algorithm.

the image—one for each detector involved in a given band. For MSS data this means that for a given band one histogram is generated for scan lines 1, 7, 13, and so on; a second is generated for lines 2, 8, 14, and so on; and so forth. These histograms are then compared in terms of their mean and median values to identify the problem detector(s). A gray-scale adjustment factor(s) can then be determined to adjust the histogram(s) for the problem lines to resemble those for the normal data lines. This adjustment factor is applied to each pixel in the problem lines and the others are not altered (Figure 7.6).

Another line-oriented noise problem sometimes encountered in digital data is *line drop*. In this situation, a number of adjacent pixels along a line (or an entire line) may contain spurious DNs. This problem is normally addressed by replacing the defective DNs with the average of the values for the pixels occurring in the lines just above and below (Figure 7.7). Alternatively, the DNs from the preceding line can simply be inserted in the defective pixels.

Figure 7.7 Line drop correction: (*a*) original image containing two line drops; (*b*) restored image resulting from averaging pixel values above and below defective line.

Random noise problems in digital data are handled quite differently than those we have discussed to this point. This type of noise is characterized by nonsystematic variations in gray levels from pixel to pixel called *bit errors*. Such noise is often referred to as being "spikey" in character, and it causes images to have a "salt and pepper" or "snowy" appearance.

Bit errors are handled by recognizing that noise values normally change much more abruptly than true image values. Thus, noise can be identified by comparing each pixel in an image with its neighbors. If the difference between a given pixel value and its surrounding values exceeds an analyst-specified threshold, the pixel is assumed to contain noise. The noisy pixel value can then be replaced by the average of its neighboring values. Moving neighborhoods or windows of 3×3 or 5×5 pixels are typically used in such procedures. Figure 7.8 illustrates the concept of a moving window comprising a 3×3-pixel neighborhood, and Figure 7.9 illustrates just one of many noise suppression algorithms using such a neighborhood. Finally, Figure 7.10 illustrates the results of applying the algorithm included in Figure 7.9 to an actual image.

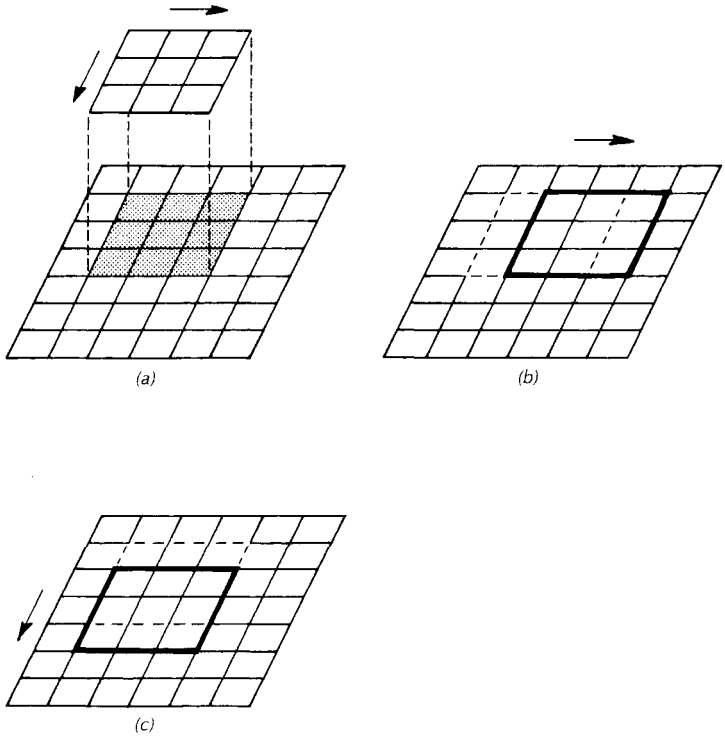

Figure 7.8 The moving window concept: (*a*) projection of 3 × 3-pixel window in image being processed; (*b*) movement of window along a line from pixel to pixel; (*c*) movement of window from line to line.

DN_1	DN_2	DN_3
DN_4	DN_5	DN_6
DN_7	DN_8	DN_9

$AVE_A = (DN_1 + DN_3 + DN_7 + DN_9)/4$
$AVE_B = (DN_2 + DN_4 + DN_6 + DN_8)/4$
$DIFF = |AVE_A - AVE_B|$
$THRESH = DIFF \times WEIGHT$
$IF: |DN_5 - AVE_A| \text{ or } |DN_5 - AVE_B| > THRESH$
$THEN: DN'_5 = AVE_B \text{ OTHERWISE } DN'_5 = DN_5$

Figure 7.9 Typical noise correction algorithm employing a 3 × 3-pixel neighborhood. *Note:* "WEIGHT" is an analyst-specified weighting factor. The lower the weight, the greater the number of pixels considered to be noise in an image.

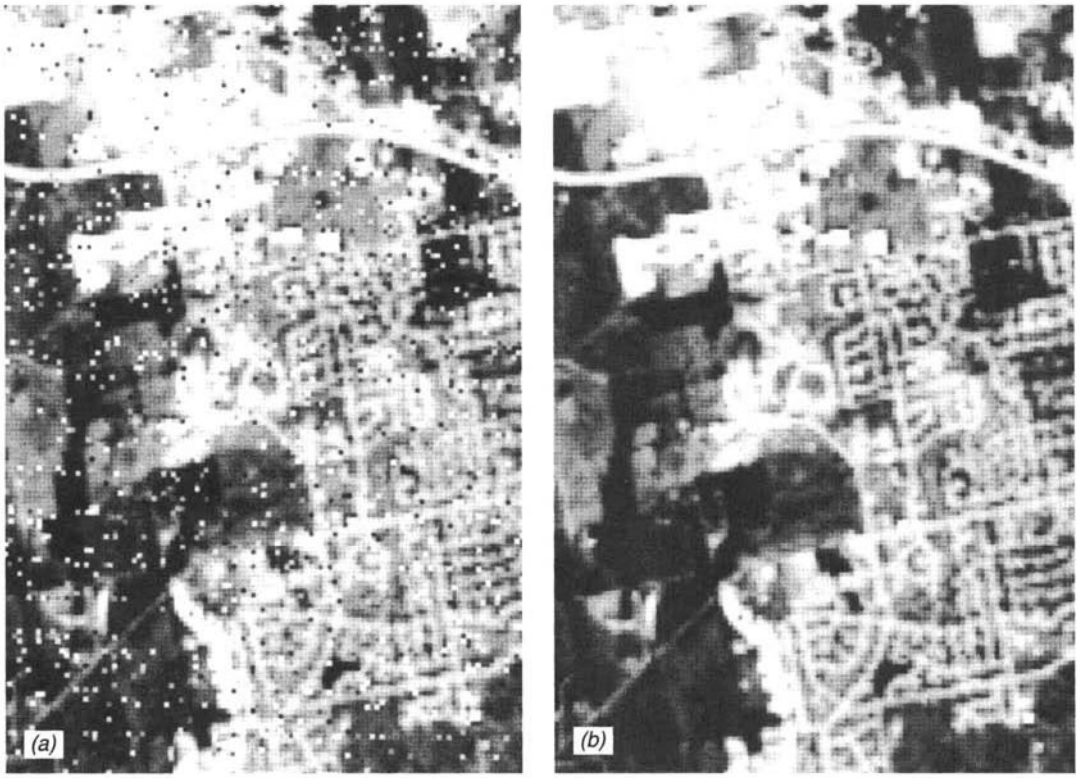

Figure 7.10 Result of applying noise reduction algorithm: (*a*) original image data with noise-induced "salt-and-pepper" appearance; (*b*) image resulting from application of algorithm shown in Figure 7.9.

7.3 IMAGE ENHANCEMENT

As previously mentioned, the goal of image enhancement is to improve the visual interpretability of an image by increasing the apparent distinction between the features in the scene. The process of visually interpreting digitally enhanced imagery attempts to optimize the complementary abilities of the human mind and the computer. The mind is excellent at interpreting spatial attributes on an image and is capable of selectively identifying obscure or subtle features. However, the eye is poor at discriminating the slight radiometric or spectral differences that may characterize such features. Computer enhancement aims to visually amplify these slight differences to make them readily observable.

The range of possible image enhancement and display options available to the image analyst is virtually limitless. Most enhancement techniques may be categorized as either point or local operations. *Point operations* modify the brightness value of each pixel in an image data set independently. *Local operations* modify the value of each pixel based on neighboring brightness values.

Either form of enhancement can be performed on single-band (monochrome) images or on the individual components of multi-image composites. The resulting images may also be recorded or displayed in black and white or in color. Choosing the appropriate enhancement(s) for any particular application is an art and often a matter of personal preference.

Enhancement operations are normally applied to image data after the appropriate restoration procedures have been performed. Noise removal, in particular, is an important precursor to most enhancements. Without it, the image interpreter is left with the prospect of analyzing enhanced noise!

Below, we discuss the most commonly applied digital enhancement techniques. Three techniques can be categorized as *contrast manipulation*, *spatial feature manipulation*, or *multi-image manipulation*. Within these broad categories, we treat the following:

1. **Contrast manipulation.** Gray-level thresholding, level slicing, and contrast stretching.
2. **Spatial feature manipulation.** Spatial filtering, edge enhancement, and Fourier analysis.
3. **Multi-image manipulation.** Multispectral band ratioing and differencing, principal components, canonical components, vegetation components, intensity–hue–saturation (IHS) color space transformations, and decorrelation stretching.

7.4 CONTRAST MANIPULATION

Gray-Level Thresholding

Gray-level thresholding is used to *segment* an input image into two classes—one for those pixels having values below an analyst-defined gray level and one for those above this value. Below, we illustrate the use of thresholding to prepare a *binary mask* for an image. Such masks are used to segment an image into two classes so that additional processing can then be applied to each class independently.

Shown in Figure 7.11*a* is a TM1 image that displays a broad range of gray levels over both land and water. Let us assume that we wish to show the brightness variations in this band in the water areas only. Because many of the gray levels for land and water overlap in this band, it would be impossible to separate these two classes using a threshold set in this band. This is not the case in the TM4 band (Figure 7.11*b*). The histogram of DNs for the TM4 image (Figure 7.11*c*) shows that water strongly absorbs the incident energy in this near-infrared band (low DNs), while the land areas are highly reflective (high DNs). A threshold set at DN = 40 permits separation of these two classes in the TM4 data. This binary classification can then be applied to the