Error Correction and Registration of Image Data

Image when recorded can contain errors in geometry and in the brightness values of the pixels. The latter are referred to:

- radiometric errors
- errors from instrumentation used
- effect of the atmosphere

Sources of Geometric Distortion

There are potentially many more sources of geometric distortion of image data than radiometric distortion and their effects are more severe. They can be rerlated to a number of factors:

- the rotation of the earth during image acquisition,
- the finite scan rate of some sensors,
- the wide field of view of some sensors,
- the curvature of the earth,
- sensor non-idealities,
- variations in platform altitude, attitude and velocity, and
- panoramic effects related to the imaging geometry.

To appreciate why geometric distortion occurs, in some cases it is necessary to envisage how an image is formed from sequential lines of image data. If one imagines that a particular sensor records L lines of N pixels each then it would be natural to form the image by laying the L lines down successively one under the other.

If the IFOV of the sensor has an aspect ratio of unity - i.e. the pixels are the same size along and across the scan then this is the same as arranging the pixels for display on a square grid, such as that shown in figure:

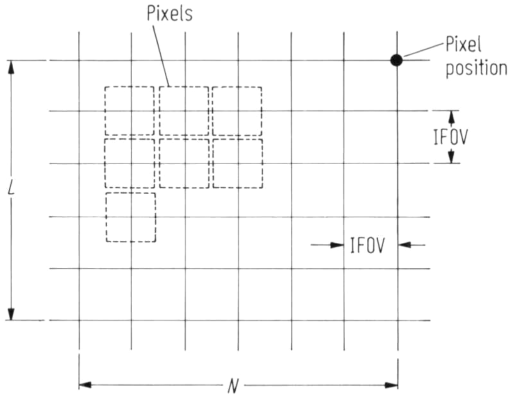

Display grid commonly used to built up an image from the digital data stream of pixels generated by a sensor

The grid intersections are the pixel positions and the spacing between those grid points is equal to the sensor's IFOV.

Earth Rotation Effects

Line scan sensors such as the Landsat MSS and TM, and the NOAA AVHRR take a finite time to acquire a frame of image data. The same is true of push broom scanners such as the SPOT HRV.

During the frame acquisition time the earth rotates from west to east so that a point imaged at the end of the frame would have been ***further*** to the west when recording started. Therefore if the lines of image data recorded were arranged for display in the manner of figure above the later lines would be erroneously displaced to the east in terms of the terrain they represent.

Instead, to give the pixels their correct positions relative to the ground it is necessary to offset the bottom of the image to the west by the amount of movement of the ground during image acquisition with all intervening lines displaced proportionately as depicted:

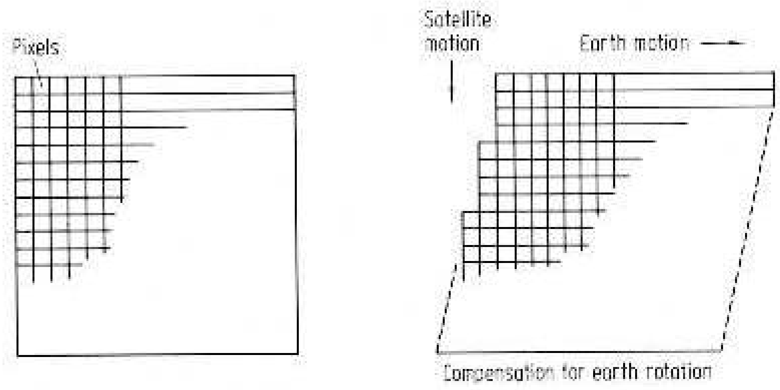

The amount by which the image has to be skewed to the west at the end of the frame depends upon the relative velocities of the satellite and earth and the length of the image frame recorded. An example is presented for Landsats 1, 2, 3.

The angular velocity of the satellite is $w_0 = 1.014$ ***mrad/s*** so that a nominal L = 185 km frame on the ground is scanned in

$$ts = L/(r_e w_0) = 28.6 \text{ s}$$

where r_e is the radius of the earth (6.37816 Mm).

The surface velocity of the earth is given by

$v_e = w_0 r_e$ ***cos ë***

where ***ë*** is latitude and w_e is the earth rotational velocity of 72.72 ***ìrad/s***. At Sydney, Australia ***ë*** = 33.80 so that

v_e = 385.4 ms^{-1}

During the frame acquisition time the surface of the earth moves to the east by

Äxe = $v_e t_s$ = 11.02 km at 33.8^0 S latitude

This represents 6% of the frame size. Since the satellite does not pass directly northsouth this movement has to be corrected by the inclination angle. At Sydney this is approximately 11^0 so that the effective sideways movement of the earth is

Äx = Äxe cos 11^0 = 10.82 km

Consequently if steps are not taken to correct an image from the first three Landsats for the effect of earth rotation then the image will contain a 6% skew distortion to the east.

Panoramic Distortion

For scanners used on spacecraft and aircraft remote sensing platforms the angular IFOV is constant. As a result the effective pixel size on the ground is larger at the extremities of the scan than at nadir, as illustrated in:

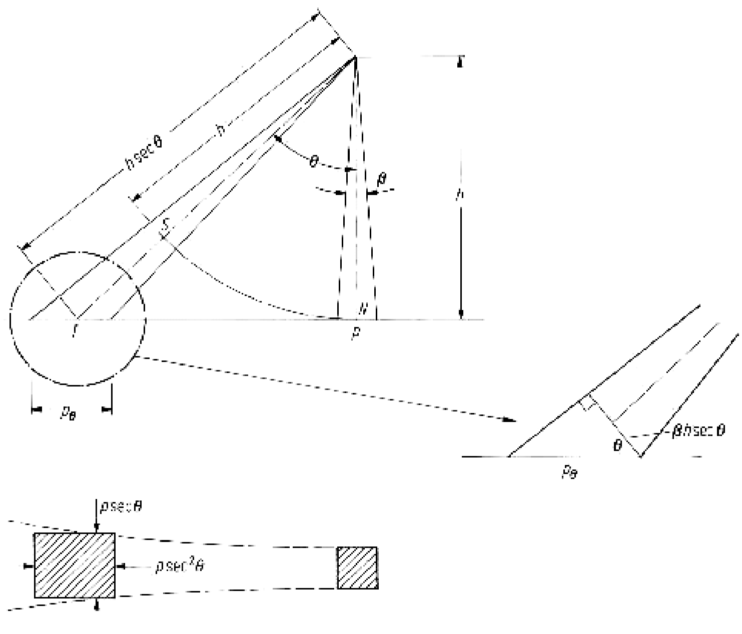

In particular, if the IFOV is â and the pixel dimension at nadir is p then its dimension in the scan direction at a scan angle of è as shown is

$$p_{è} = â\ h\ sec^2\ è = p\ sec^2\ è$$

where h is altitude. For small values of è these effects are negligible. For example, for Landsats 4 and 5 the largest value of è is approximately 7.5^0 so that $p_{è} = 1.02$ p. However for systems with larger fields of view the effect can be quite severe. For an aircraft scanner with FOV = 80^0 the distortion in pixel along the scan line is $p_{è} = 1.70\ p$ - i. e. the region on the ground measured at the extremities of the scan is 70% larger laterally than the region sensed at nadir.

There is a second distortion introduced with wide field of view systems and that relates to pixel positions across the scan line. The scanner records pixels at constant angular increments and these are displayed on a grid of uniform centres.

However the spacings of the effective pixels on the ground increase with scan angle. For example if the pixels are recorded at an angular separation equal to the IFOV of the sensor then at nadir the pixels centres are spaced p apart. At a scan angle è the pixel centres will be spaced $p\ sec^2\ è$ apart.

Thus by placing the pixels on a uniform display grid the image will suffer an across track compression. Again the effect for small angular field of view systems will be negligible in terms of the relative spacings of adjacent pixels. However when the effect is compounded to determine the location of a pixel at the swath edge relative to nadir the error can be significant.

These panoramic effects lead to an interesting distortion in the geometry of large field of view systems. To see this consider the uniform mesh shown in:

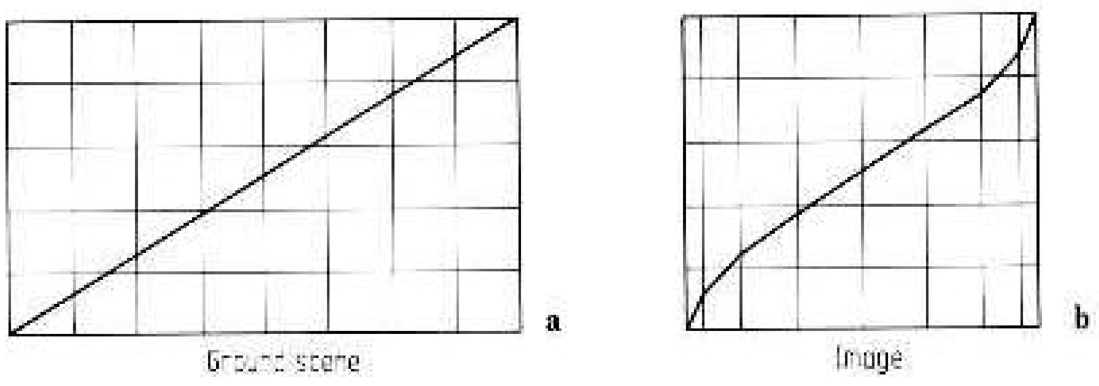

Suppose this represents a region on the ground being imaged. For simplicity the cells in the grid could be considered to be features on the ground. Because of the compression in the image data caused by displaying equal-sized pixels on a uniform grid as discussed in the foregoing, the uniform mesh will appear as shown in (b). Image pixels are recorded with a constant IFOV and at a constant angular sampling rate.

The number of pixels recorded therefore over the outer grid cells in the along scan direction will be smaller than over those near nadir. In the along track direction there is no variation of pixel spacing or density with scan angle as this is established by the forward motion of the platform. Rather pixels near the swath edges will contain information in common owing to the overlapping IFOV.

Linear features such as roads at an angle to the scan direction will appear **bent** in the displayed image data because of the along scan compression effect. Owing to the change in shape caused, the distortion is frequently referred to as **S-bend** distortion and can be a common problem with aircraft line scanners. Clearly, not only linear features are affected; rather the whole image detail near the swath edges is distorted in this manner.

Earth Curvature

Aircraft scanning systems, because of their low altitude (and thus the small absolute swath width of their image data), are not affected by earth curvature. Neither are space systems such as Landsat and SPOT, again because of the narrowness of their swaths.

However wide swath width spaceborne imaging systems are affected. For NOAA with a swath width of 2700 km and an altitude of 833 km it can be shown that the deviation of the earth's surface from a plane amounts to about 2.3 % over the swath, which seems insignificant.

However it is the inclination of the earth's surface over the swath that causes the greater effect. At the edges of the swath the area of the earth's surface viewed at a given angular IFOV is larger than if the curvature of the earth is ignored. The increase in pixel size can be computed by reference to the geometry of figure below.

The pixel dimension in the across track direction normal to the direction of the sensor is *â [h +* r_c *(1 - cos ö)] sec è* as shown. The geometry then shows that the effective pixel size on the inclined earth's surface is:

$$p_c = \beta[h + r_e(1 - \cos\phi)]\sec\theta\sec(\theta + \phi)$$

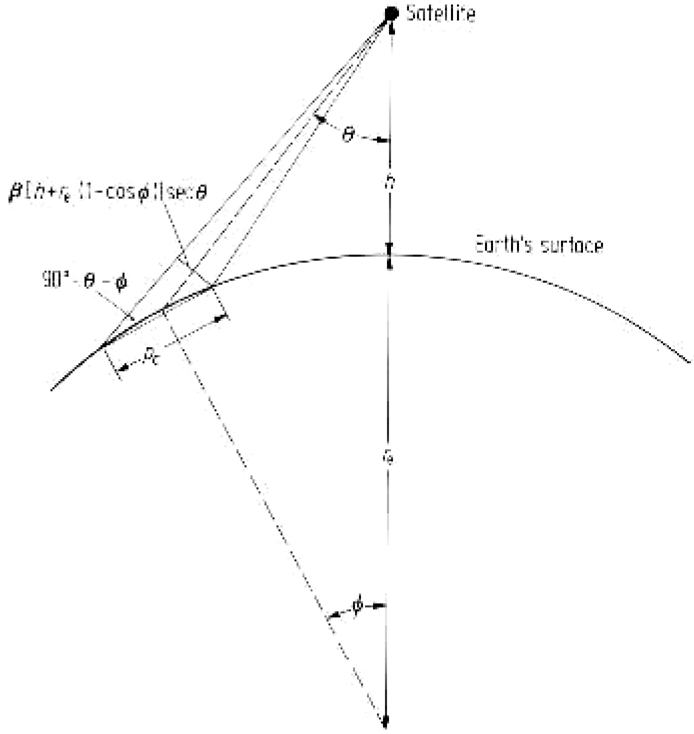

where *â h* is the pixel size at nadir and ö is the angle subtended at the centre of the earth. Using the NOAA satellite as an example è = 54' at the edge of the swath and ö = 12^0. This shows that the effective pixel size in the along scan direction is 2.89 times larger than that at nadir when earth curvature is ignored, but is 4.94 times that at nadir when the effect of earth curvature is included.

This demonstrates that earth curvature introduces a significant additional compressive distortion in the image data acquired by satellites such as NOAA when an image is constructed on a uniform grid. The effect of earth curvature in the along track direction is negligible.

Scan Time Skew

Mechanical line scanners such as the Landsat MSS and TM require a finite time to scan across the swath. During this time the satellite is moving forward leading to a skewing in the along track direction. As an illustration of the magnitude of the effect, the time require to record one MSS scan line of data is 33 ms. During this time the satellite travels forward by 213 m at its equivalent ground velocity of 6.467 km/s. As a result the end of the scan line is advanced by this amount compared with its start.

Variations in Platform Altitude, Velocity and Attitude

Variations in the elevation or altitude of a remote sensing platform lead to a scale change at constant angular IFOV and field of view; the effect is illustrated in:

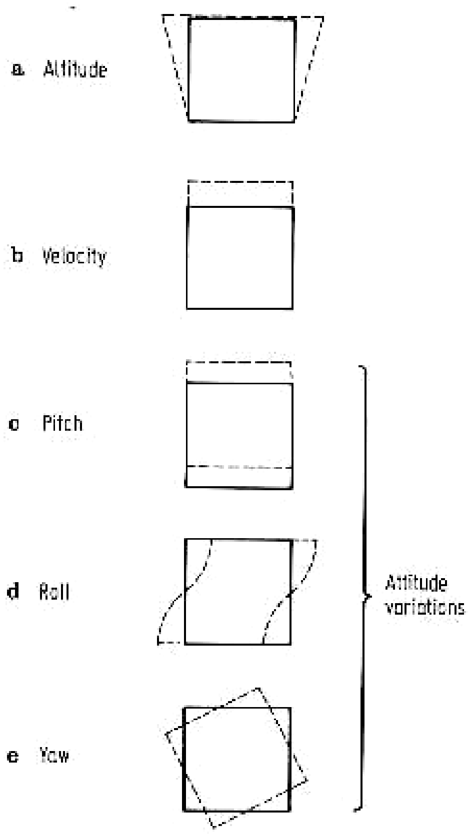

This effect is for an increase in altitude with travel at a rate that is slow compared with a frame acquisition time. Similarly, if the platform forward velocity changes, a scale change occurs in the long track direction. For a satellite platform, orbit velocity variations can result from orbit eccentricity and the non-sphericity of the earth.

Platform attitude changes can be resolved into yaw, pitch and roll during forward travel. These lead to image rotation, along track and across track displacement.

While these variations can be described mathematically, at least in principle, a knowledge of the platform ephemeris is required to enable their magnitudes to be computed. In the case of satellite platforms ephemeris information is often telemetered to ground receiving stations. For the Landsat system this is used to apply corrections before the data is distributed.

Attitude variations in aircraft remote sensing systems can potentially be quite significant owing to the effects of **atmospheric turbulence**. These can occur over a short time, leading to localised distortions in aircraft scanner images. Frequently aircraft roll is compensated for in the data stream. This is made possible by having a data window that defines the swath width; this is made smaller than the complete scan of data over the sensor field of view. A gyro mounted on the sensor is then used to move the position of the data window along the total scan line as the aircraft rolls. Pitch and yaw are generally not corrected unless the sensor is mounted on a three axis stabilized platform.

Aspect Ratio Distortion

The aspect ratio of an image (that is, its scale vertically compared with its scale horizontally) can be distorted by mechanisms that lead to overlapping IFOV's. The most notable example of this occurs with the Landsat multispectral scanner. Samples are taken across a scan line too quickly compared with the IFOV. This leads to pixels having 56 metre centres but sampled with an IFOV of 79 m.

Consequently the effective pixel size is 79 m x 56 m and thus is not square. As a result if the pixels recorded by the multispectral scanner are displayed on the square grid, the image will be too wide for its height when related to the corresponding region on the ground. The magnitude of the distortion is 79/56 = 1.411 so that this is quite a severe error and must be corrected for most applications.

A similar distortion can occur with aircraft scanners if the velocity of the aircraft is not matched to the scanning rate of the sensor. Either underscanning or overscanning can occur leading to distortion in the alongtrack scale of the image.

Sensor Scan Nonlinearities

Line scanners that make use of rotating mirrors, such as the NOAA AVHRR and aircraft scanners, have a scan rate across the swath that is constant, to the extent that the scan motor speed is constant.

Systems that use an oscillating mirror however, such as the Landsat multispectral scanner, incur some nonlinearity in scanning near the swath edges owing to the need for the mirror to slow down and change directions:

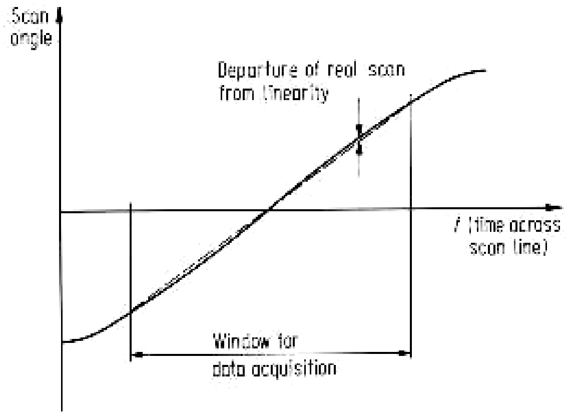

Correction of Geometric Distortion

There are two techniques that can be used to correct the various types of geometric distortion present in digital image data. One is ***to model*** the nature and magnitude of the sources of distortion and use these models to establish correction formulae. This technique is effective when the types of distortion are well characterized, such as that caused by earth rotation.

The second approach depends upon establishing **mathematical** relationships between the addresses of pixels in an image and the corresponding coordinates of those points on the ground (via a map). These relationships can be used to correct the image geometry irrespective of the analyst's knowledge of the source and type of distortion.

This procedure will be treated first since it is the most commonly used and, as a technique, is independent of the platform used for data acquisition. It should be noted that each band of image data has to be corrected. However since it can often be assumed that the bands are well registered to each other, steps taken to correct one band in an image, can be used on all remaining bands.

Use of Mapping Polynomials for Image Correction

An assumption that is made in this procedure is that there is available a map of the region corresponding to the image, that is correct geometrically. We then define two Cartesian coordinate systems as shown below. One describes the location of points in the map (x, y) and the other coordinate system defines the location of pixels in the image (u, v). Now suppose that the two coordinate systems can be related via a pair of mapping functions f and g so that:

$$u = f(x, y), \quad v = g(x, y)$$

If these functions are known then we could locate a point in the image knowing its position on the map. In principle, the reverse is also true. With this ability we could build up a geometrically correct version of the image in the following manner. First we define a grid over the map to act as the grid of pixel centres in the corrected image. This grid is parallel to the map coordinate grid itself, described by latitudes and longitudes, UTM coordinates and so on.

We will refer to this grid as the display grid; by definition this is geometrically correct. We then move over the display grid pixel centre by pixel centre and use the mapping functions above to find the corresponding pixel in the image for each display grid position. Those pixels are then placed on the display grid. At the conclusion of the process we have a geometrically correct image built up on the display grid utilizing the original image as a source of pixels.

While the process is a straightforward one there are some practical difficulties that must be addressed. First we do not know the explicit form of the mapping functions of *u*, *v*. Secondly, even if we did, they may not point exactly to a pixel in the image corresponding to a display grid location; instead some form of interpolation may be required.

Mapping Polynomials and Ground Control Points

Since explicit forms for the mapping functions are not known they are generally chosen as simple polynomials of first, second or third degree. For example, in the case of second degree (or order)

$$u = a_0 + a_1 x + a_2 y + a_3 xy + a_4 x^2 + a_5 y^2$$
$$v = b_0 + b_1 x + b_2 y + b_3 xy + b_4 x^2 + b_5 y^2$$

Sometimes orders higher than third are used but care must be taken to avoid the introduction of worse errors than those to be corrected.

If the coefficients a_i and b_i were known then the mapping polynomials could be used to relate any point in the map to its corresponding point in the image as in the foregoing discussion. At present however these coefficients are unknown. Values can be estimated by identifying sets of features on the map that can also be identified on the image.

These features, often referred to as ***ground control points*** (GCP's) are welldefined and spatially small and could be road intersections, airport runway intersections, bends in rivers, prominent coastline features and the like. Enough of these are chosen (as pairs on the map and image) so that the polynomial coefficients can be estimated by substitution into the mapping polynomials to yield sets of equations in those unknowns. The minimum number required for second order polynomial mapping is six Likewise a minimum of three is required for first order mapping and ten for third order mapping.

In practice however significantly more than these are chosen and the coefficients are evaluated using least squares estimation. In this manner any control points that contain significant positional errors either on the map or in the image will not have an undue influence on the polynomial coefficients.

Resampling

Having determined the mapping polynomials explicitly by use of the ground control points the next step is to find points in the image corresponding to each location in the pixel grid previously defined over the map. The spacing of that grid is chosen according to the pixel size required in the corrected image and need not be the same as that in the original geometrically distorted version. For the moment suppose that the points located in the image correspond exactly to image pixel centres. Then those pixels are simply transferred to the appropriate locations on the display grid to build up the rectified image. This is the case as in figure:

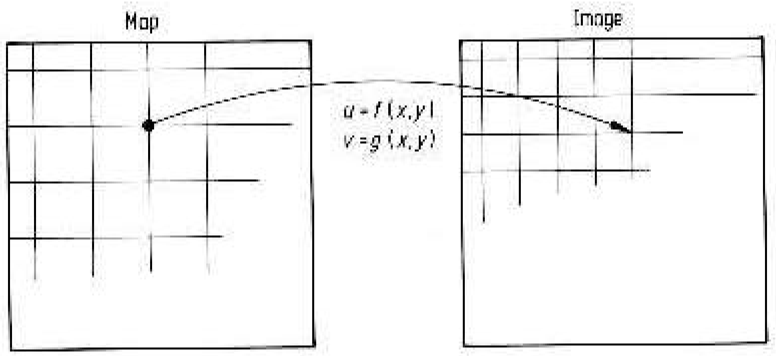

Interpolation

As is to be expected, grid centres from the map-registered pixel grid will not usually project to exact pixel centre locations in the image, and some decision has to be made therefore about what pixel brightness value should be chosen for placement on the new grid. Three techniques can be used for this purpose.

Nearest neighbour resampling simply chooses the actual pixel that has its centre nearest the point located in the image, as depicted in figure below (a). This pixel is then transferred to the corresponding display grid location. This is the preferred technique if the new image is to be classified since it then consists of the original pixel brightnesses, simply rearranged in position to give a correct image geometry.

Bilinear interpolation uses three linear interpolations over the four pixels that surround the point found in the image corresponding to a given display grid position. The process is illustrated in figure below (b). Two linear interpolations are performed along the scan lines to find the interpolants.

Cubic convolution interpolation uses the surrounding sixteen pixels. Cubic polynomials are fitted along the four lines of four pixels surrounding the point in the image, as depicted in figure below (c) to form four interpolants. A fifth cubic polynomial is then fitted through these to synthesise a brightness value for the corresponding location in the display grid.

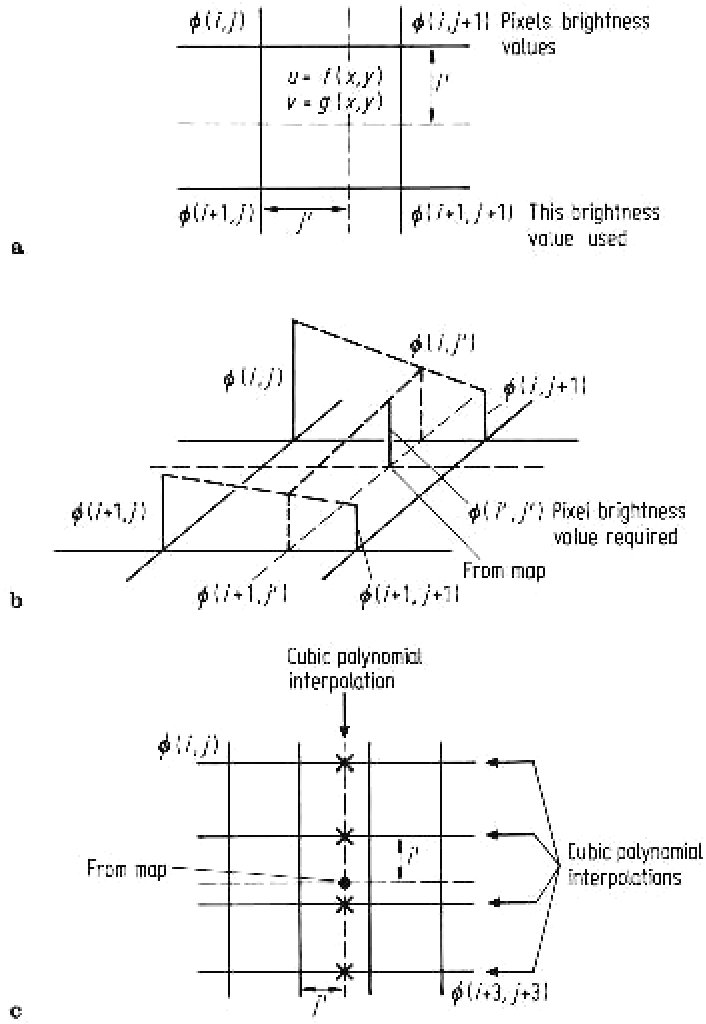

The actual form of polynomial that is used for the interpolation is derived from considerations in sampling theory and issues concerned with constructing a continuous function (i.e. interpolating) from a set of samples.

Cubic convolution interpolation, or resampling, yields an image product that is generally smooth in appearance and is often used if the final product is to be treated by photointerpretation. However since it gives pixels on the display grid, with brightnesses that are interpolated from the original data, it is not recommended if classification is to follow since the new brightness values may be slightly different to the actual radiance values detected by the satellite sensors.

Choice of Control Points

Enough well defined control point pairs must be chosen in rectifying an image to ensure that accurate mapping polynomials are generated. However care must also be given to the locations of the points. A general rule is that there should be a distribution of control points around the edges of the image to be corrected with a scattering of points over the body of the image. This is necessary to ensure that the mapping polynomials are well-behaved over the image.

This concept can be illustrated by considering an example from curve fitting. While the nature of the problem is different the undesirable effects that can be generated are similar. Note that as the order is higher the curves pass closer to the points. However if it is presumed that the data would have continued for larger values of x with much the same trend as apparent in the points plotted then clearly the linear fit will extrapolate moderately acceptably. In contrast the cubic curve can deviate markedly from the trend when used as an extrapolator. This is essentially true in geometric correction of image data: while the higher order polynomials will be accurate in the vicinity of the control points themselves, they can lead to significant errors, and thus image distortions, for regions of images outside the range of the control points.

Mathematical Modelling

If a particular distortion in image geometry can be represented mathematically then the mapping functions can be specified explicitly. This obviates the need to choose arbitary polynomials and to use control points to determine the polynomial coefficients. In this section some of the more common distortions are treated from this point of view.

However rather than commence with expressions that relate image coordinates (u, v) to map coordinates (x, y) it is probably simpler conceptually to model what the true (map) positions of pixels should be given their positions in an image. This expression can then be inverted if required to allow the image to be resampled on to the map grid.

Aspect Ratio Correction

The easiest source of distortion to model is that caused by the 56 m equivalent ground spacing of the 79m *x* 79m equivalent pixels in the Landsat multispectral scanner. This leads to an image that is too wide for its height by a factor of 79/56 = 1.411.

Consequently to produce a geometrically correct image either the vertical dimension has to be expanded by this amount or the horizontal dimension must be compressed. We will consider the former. This requires that the pixel axis horizontally be left unchanged (i.e. $x=u$), but that the axis vertically be scaled (i.e. $y = 1.411\ v$). This can be expressed conveniently in matrix notation as

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1.411 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

One way of implementing this corection would be to add extra lines of pixel data to expand the vertical scale. This could be done by duplicating four lines in every ten. Alternatively, and more precisely, the equation above can be inverted to give

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0.709 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

Thus, a display grid is defined over the map is used to find the corresponding location in the image (u, v). The interpolation techniques are then used to generate brightness values for the display grid pixels.

Earth Rotation Skew Correction

To correct for the effect of earth rotation it is necessary to implement a shift of pixels to the left that is dependent upon the particular line of pixels measured with respect to the top of the image. Their line addresses as such (v) are not affected. These corrections are implemented by

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} 1 & a \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

with *á* having value of local coordinates. Again this can be implemented in an approximate sense by making one pixel shift to the left every *N* lines of image data measured down from the top.

Image Orientation to North-South

Although not strictly a geometric distortion it is an inconvenience to have an image that is corrected for most major effects but is not oriented vertically in a north-south direction. It will be recalled for example that the Landsat orbits in particular are inclined to the north-south line by about 9'. (This of course is different with different latitudes). To rotate an image by an angle î in the counter or anticlockwise direction (as required in the case of Landsat):

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \cos x & \sin x \\ -\sin x & \cos x \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

so that

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \cos x & -\sin x \\ \sin x & \cos x \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

Correction of Panoramic Effects

Let's make note of the pixel positional error that results from scanning with a fixed IFOV at a constant angular rate. In terms of map and image coordinates, the distortion can be described by

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \tan q / q & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

where *è* is the instantaneous scan angle, which in turn can be related to *x* or *u*, viz. *x = h tan è, u = h è,* where h is altitude. Consequently resampling can be carried out according to

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} q \cot q & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} h \tan^{-1}\left(\frac{x}{h}\right) / x & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

Image Registration

Georeferencing and Geocoding

Using the correction techniques an image can be registered to a map coordinate system and therefore have its pixels addressable in terms of map coordinates (eastings and northings, or latitudes and longitudes) rather than pixel and line numbers. Other spatial data types, such as geophysical measurements, image data from other sensors and the like, can be registered similarly to the map thus creating a georeferenced integrated spatial data base of the type used in a geographic information system.

Expressing image pixel addresses in terms of a map coordinate base is often referred to as geocoding; ultimately it would be anticipated that remote sensing image data could be purchased according to bounds expressed in map coordinates rather than in scenes or frames.

Image to Image Registration

Many applications of remote sensing image data require two or more scenes of the same geographical region, acquired at different dates, to be processed together. Such a situation arises for example when changes are of interest, in which case registered images allow a pixel by pixel comparison to be made.

Two images can be registered to each other by registering each to a map coordinate base separately, in the manner demonstrated in the previous section. Alternatively, and particularly if georeferencing is not important, one image can be chosen as a master to which the other, known as the slave, is to be registered. The coordinates (x, y) are now the pixel coordinates in the master image rather than the map coordinates, and (u, v) are the coordinates of the image to be registered (i.e. the slave). An advantage in image to image registration is that only one registration step is required in comparison to two if both are taken back to a map base. Furthermore an artifice known as a sequential similarity detection algorithm can be used to assist in accurate co-location of control point pairs.

Sequential Similarity Detection Algorithm

A correlation procedure is of value in locating the position of a control point in the master image having identified it in the slave. Known as a sequential similarity detection algorithm (SSDA), the technique has several variations, and a specific implementation is considered here to illustrate the nature of the method.

Suppose a control point has been chosen in the slave image and it is necessary to determine its counterpart in the master image. In principle a rectangular sample of pixels surrounding the control point in the slave image can be extracted as a window to be correlated with the master image. Because of the spatial properties of the pair of images near the control points a high correlation should occur when the slave window is located over its exact counterpart region in the master, and thus the master location of the control point is identified.

Obviously it is not necessary to move the slave window over the complete master image since the user knows approximately where the control point should occur in the master. Consequently it is only necessary to specify a search region in the neighbourhood of the approximate location. Software systems that provide this option allow the user to choose both the size of the window of pixels from the slave image control point neighbourhood and the size of the search region in the master image over which the window of slave pixels is moved to detect an acceptable correlation.

The correlation measure used need not be sophisticated. Indeed a simple similarity check that can be used is to compute the sum of the absolute differences of the slave and master image pixel brightnesses over the window, for each possible location of the window in the search region. The location that gives the smallest absolute difference defines the control point position as that pixel at the current centre of the window. Obviously the sensitivity of the method will be reduced if there is a large average difference in brightness between the two images - such as that owing to seasonal variations. A refinement therefore is to compute the summed absolute difference of the pixel brightnesses relative to their respective means in the search window.

Clearly the use of techniques such as these to locate control points depends upon there not being major changes of an uncorrelated nature between the scenes in the vicinity of a control point being tested. For example a vegetation flush due to rainfall in part of the search window can lead to an erroneous location. Nevertheless with a judicious choice of window size and search region, measures such as SSDA can give very effective guidance to the user, especially when available on an interactive image processing facility.

Example of Image to Image Registration

To illustrate image to image registration, but more particularly to see clearly the effect of control point distribution and the significance of the order of the mapping polynomials to be used for registration, two segments of Landsat MSS infrared image data in the northern suburbs of Sydney were chosen. One was acquired on December 29, 1979 and was used as the master. The other was acquired on

December 14, 1980 and was used as the slave image. These are shown in figure below wherein careful inspection shows the differences in image geometry.

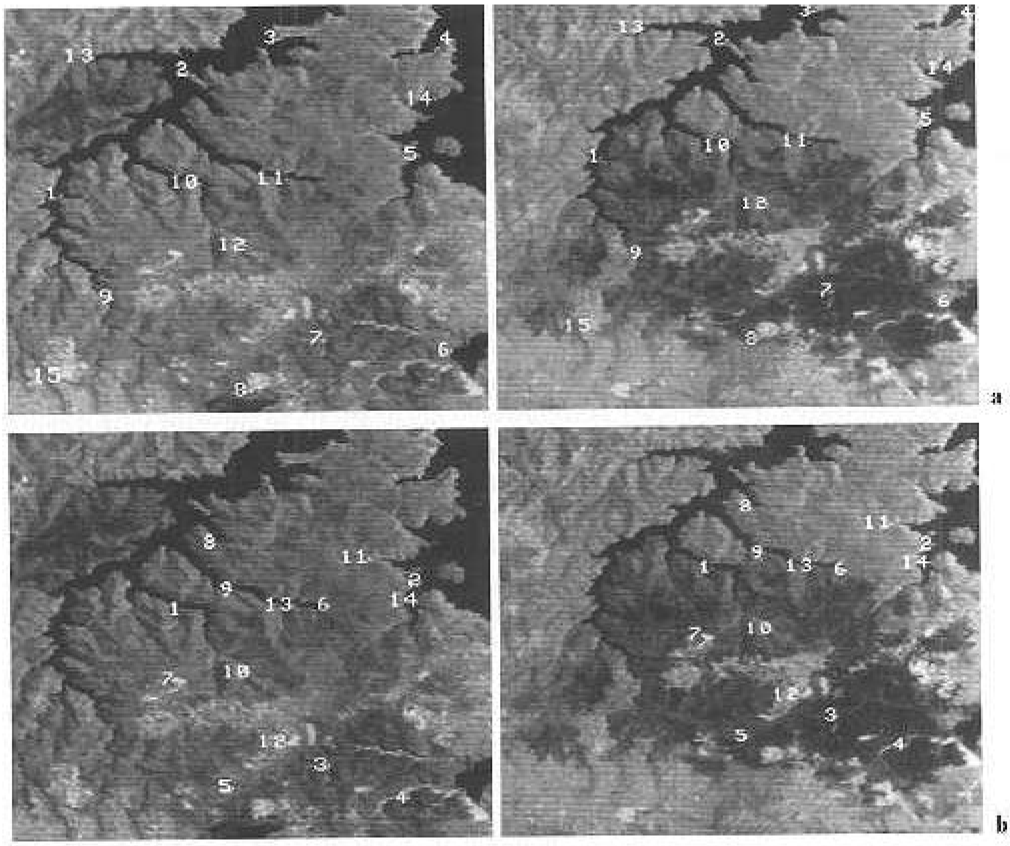

Two sets of control points were chosen. In one the points were distributed as nearly as possible in a uniform manner around the edge of the image segment as shown in a, with some points located across the centre of the image. This set would be expected to give a reasonable registration of the images.

The second set of control points was chosen injudiciously, closely grouped around one particular region, to illustrate the resampling errors that can occur. These are shown in b. In both cases the control point pairs were co-located with the assistance of a sequential similarity detection algorithm. This worked well particularly for those control points around the coastal and river regions where the similarity between the images is unmistakable.

To minimise tidal influences on the location of control points, those on water boundaries were chosen as near as possible to be on headlands, and certainly were never chosen at the end of inlets.

For both sets of control points third degree mapping polynomials were used along with cubic convolution resampling. As expected the first set of points led to an acceptable registration of the images whereas the second set gave a good registration in the immediate neighbourhood of the points but beyond them produced gross distortion.

The adequacy of the registration process can be assessed visually if the master and resampled slave images can be superimposed in different colours using a colour display system:

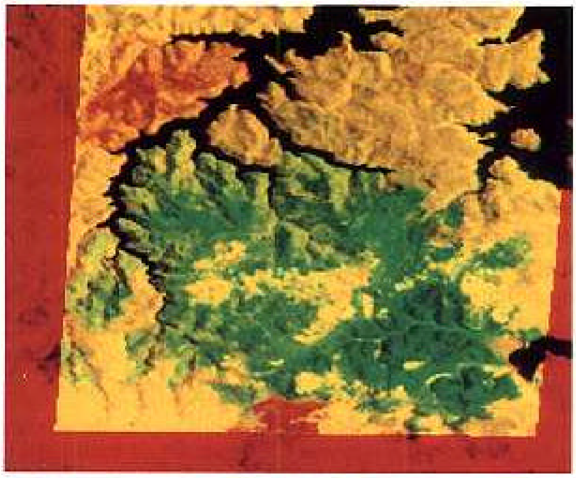

a) registration of the 1980 image (green) with 1979 image (red) using the control points of figure before a, with third order mapping polynomials

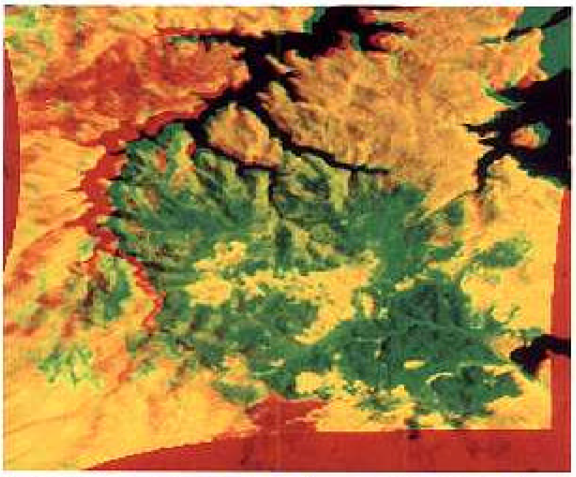

b) third order mapping of 1980 image (green) to 1979 image (red) using the control points of figure before b.

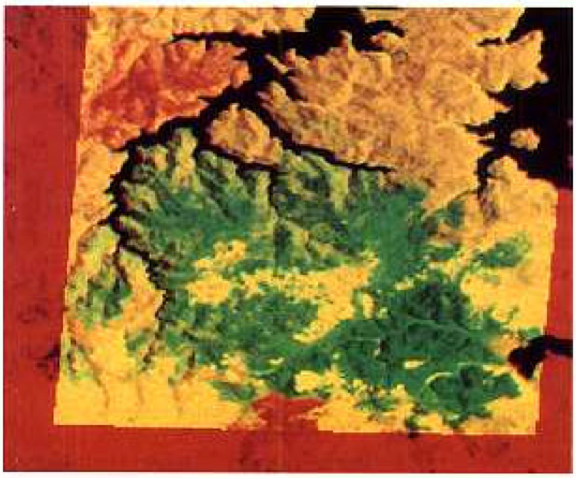

c) as for b) but using the first order mapping polynomials

Figures a and b show the master image in red with the resampled slave image superimposed in green. Where good registration has been achieved the resultant is yellow (with the exception of regions of gross dissimilarity in pixel brightness - in this case associated with fire burns). However misregistration shows quite graphically by a red-green separation. This is particularly noticeable in b, where the poor extrapolation obtained with third order mapping is demonstrated.

The exercise using the poor set of control points b was repeated. However this time only first order mapping polynomials were used. While these obviously will not remove non-linear differences between the images and will give poorer matches at the control points themselves, they are well behaved in extrapolation beyond the vicinity of the control points and lead to an acceptable registration as shown in c.