Field Techniques Manual:
GIS, GPS and Remote Sensing

- ## Section C: Techniques

Chapter 9: Geocorrection and photogrammetry

9 Geocorrection and Photogrammetry

Aerial photographs are, arguably, the most widely used form of remote sensing. Airphotos have been used in the production and update of almost all topographic maps since the 1950s in almost every country in the world. It is therefore highly likely that photographs already exist of the area you intend to visit, wherever that happens to be. Unfortunately, you may not be able to access airphotos of your study area, most commonly because of military sensitivity, but perhaps due to other factors, such as local government policy, copyright issues and depleted airphoto archives. Further problems encountered by the editorial team on expeditions in developing countries include a lack of chemicals to develop prints and termite damage to airphoto archives. The scale, quality and age (date acquired) of the photography could also be very variable, again dependent on why the photographs were taken.

Measurement from aerial photographs is called *photogrammetry*. The term is derived from three Greek root words and literally means *light-writing-measurement*. The science of photogrammetry is solely concerned with collecting precise and accurate measurements from photographs. Photogrammetry is one of the more complex GISci subjects, so this chapter will focus on practical aspects that are of use in field work. For more in-depth coverage of photogrammetry theory and techniques, the reader is referred to Avery & Berlin (1992), Lillesand & Kiefer (2000), Mikhail *et al.* (2001) and Wolf & Dewitt (2000). A review of software with photogrammetric capabilities is given in the GISci software chapter.

Most aerial photographs are acquired using specialised mapping cameras and are acquired with the camera pointing vertically down from the aircraft. Other photography may exist, notably oblique photos, such as ‘tourist snaps’ from aircraft windows or non-vertical specialist mapping-camera shots, but these are not as useful for mapping as vertical photography. An example of an oblique database is that held by English Heritage – these were primarily acquired for archaeological investigations. The reason why most photographs are acquired vertically and with specialised cameras is so that measurements of objects and areas on the Earth’s surface can be taken with the minimum of calculation and correction for distortion. Distortion is still present in the photography - as it is in all photography, terrestrial or airborne - but the vertical viewpoint reduces it to manageable proportions.

Aerial photography is typically acquired with significant overlap between images. This is done to produce views of the Earth’s surface from two different perspectives to allow ‘stereo-viewing’. When one image is viewed with your left eye and the other with your right eye, your brain integrates the images and the effect is a 3-D view of the area.

9.1 Data from space-borne sensors

A number of satellite systems exist that can produce imagery suitable for stereo-viewing. Notable systems are Corona (1960s-1970s US spy satellite images, many now declassified), SPOT (French satellite series, collecting data since the mid 1980s), IKONOS (since 1999, 1 m resolution commercial satellite series from US) and ASTER (since 2000, Japan/NASA collaboration). NASA have an archive of colour photographs from the Metric Camera (10-20 m pixels) taken during orbits of the Space Shuttle. ASTER and SPOT data can also be processed digitally to derive elevation information (DEMs) using photogrammetry; however this is a complicated and time-consuming process that requires access to expensive software for the best results. A much cheaper and quicker means of getting DEMs of your study area is to use data collected by the USA's Shuttle Radar Topographic Mission (SRTM). The SRTM DEM coverage is world-wide and can yield contours at 8-10 m intervals with 90 m pixels (30 m over the USA – hopefully this more detailed coverage will eventually be available beyond North America). One of the best things about the SRTM data is that it can be downloaded free via the internet (either from the GLCF site or from the US Geological Survey the latter giving better accuracy and being easier to mosaic together – see the Appendix and Weblinks section for contact details). Furthermore, software to process the SRTM DEMs, producing 3-D views and virtual fly-overs is also freely available via the Internet. A tutorial detailing how to download SRTM data from the GLCF website, plus how to download the *3-DEM* software and use it to generate DEMs from the SRTM data, is included in the CD that accompanies this handbook.

Many satellite data suppliers provide images that are geo-located, but the accuracy of this positional information should always be verified independently - the author has personal experience of corrected radar data with positional errors of over 300 m - and other sources have found positional errors of up to 1 km with ASTER imagery. Landsat ETM+ imagery, derived from the ESDI site (see appendix for web link and further information), is generally accurately located and can actually be used to correct similar data (such as ASTER, satellite radar, etc) to within 50 m, using image-to-image geocorrection with the Landsat image treated as the 'true position' map reference dataset. Although focused on aerial photography, this chapter provides an introduction to the techniques that you could use to make precise measurements, or generate DEMs, from all types of stereoscopic imagery, be they airphotos, spy satellite photos or satellite scanner images.

9.2 Photogrammetry

9.2.1 Perspective distortions

As mentioned earlier, all photography contains distortions. Unfortunately, these distortions, which are caused by the camera lens, the topography of the imaged surface (the Earth's surface) and the single-point instantaneous collection nature of photographs, mean that aerial photos cannot be used directly as maps. Attempts to derive distances, areas or angular relationships directly from photography will result in serious errors. Photogrammetric analysis attempts to account for these distortions to allow distances, areas and angles to be derived accurately and precisely. Most of the distortion present in photography is the result of the perspective viewpoint from which the photograph was taken. This perspective is essential, however, for stereo-viewing and is necessary for calculating elevation from photography. The perspective distortion in a vertical aerial photograph increases away from the centre of the photograph. This is because the centre of the photo (correctly called the principal point) is the only truly vertical part of the image. Figure 9-1, Figure 9-2 and Figure 9-3 illustrate this.

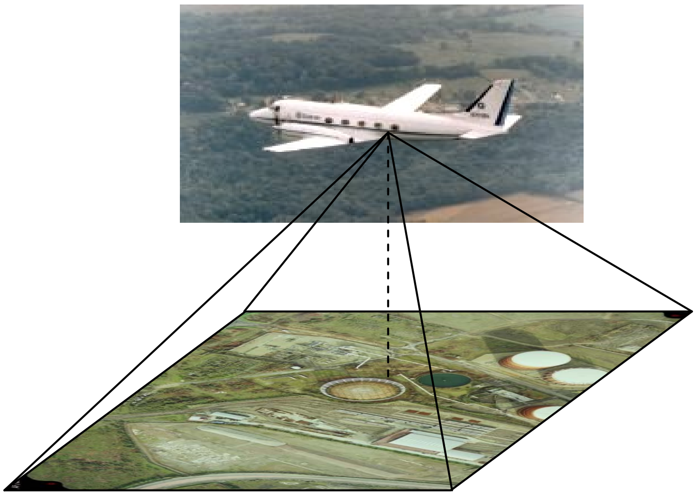

Figure 9-1 Aerial photo perspective and principal point (the only truly vertical point on the photograph is shown by the dashed line).

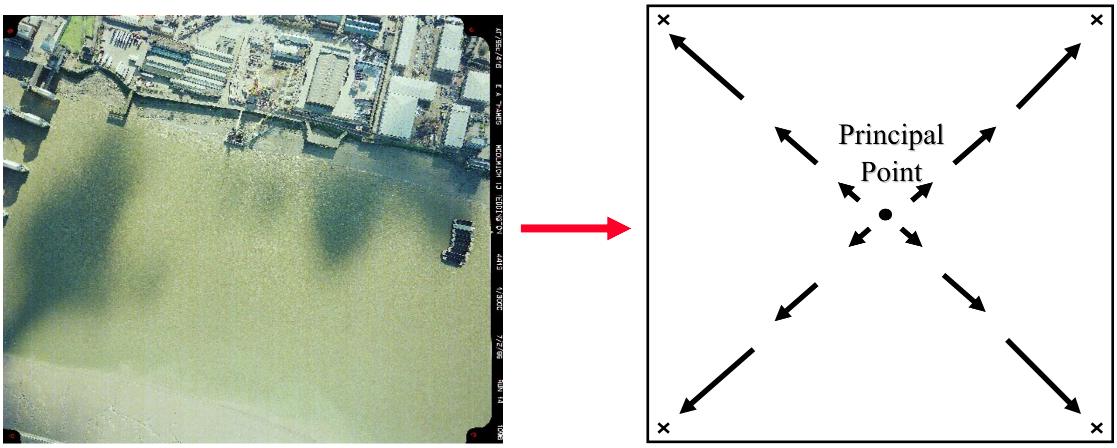

Figure 9-2 Radial distortion in a photograph caused by photographic perspective. Crosses in the corners are the fiducial marks (see Figure 9-7).

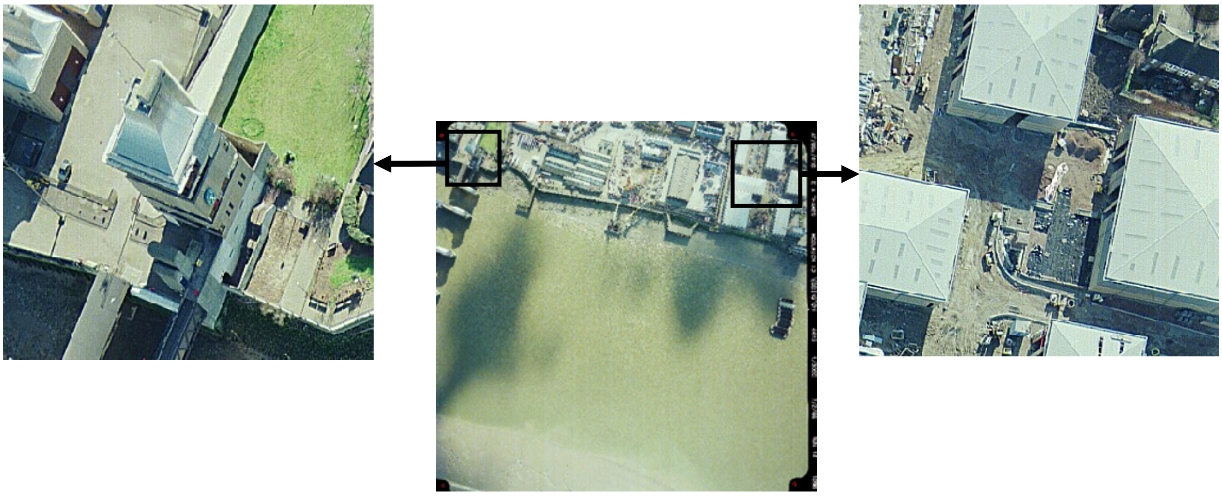

Figure 9-3 Radial distortion effects on buildings. Note how the buildings appear to lean away from the principal point (the photos are of the Thames Barrier, the river Thames and surrounding areas)

The effect of perspective distortion is that vertical buildings, hills and trees appear to 'lay-over' on their side as shown in Figure 9-3. At the principal point of a vertical air photo, the only part of a building that is visible is the roof. The roofline is actually what is used to define a building on a planimetrically correct map. The further away from the principal point a feature is, the larger the 'layover' effect gets as the photo becomes progressively closer to horizontal (Figure 9-2, Figure 9-3). An interesting result of this is that tall buildings, trees and other structures will appear to radiate away from the principal point of a vertical air photo.

The effects of terrain and radial distortion are illustrated in Plate 14, using a home experiment set-up. Three of the glasses in Plate 14 (upper left, lower left and lower right – 3(b)) are equally spaced, with 20 cm between their centres and the centre glass. The fourth glass (upper right – 3(b)) is 15 cm from the centre glass. In the vertical photo, the glasses appear to radiate outwards from the centre of the image, with the closer glass apparently less tilted and the top of the raised glass apparently further away. All these effects are systematic and can be used to help measure features and objects in three dimensions from vertical photography.

9.2.2 Parallax explained

On successive photographs, different sides of a particular feature or building may be visible. A tall feature may also 'layover' in a completely different direction on two successive photos. This effect is due to the change in location of the aeroplane between photo exposures. The apparent change in the relative positions of the bottom and top of a non-horizontal feature is called parallax. Parallax between two images is illustrated in Figure 9-4.

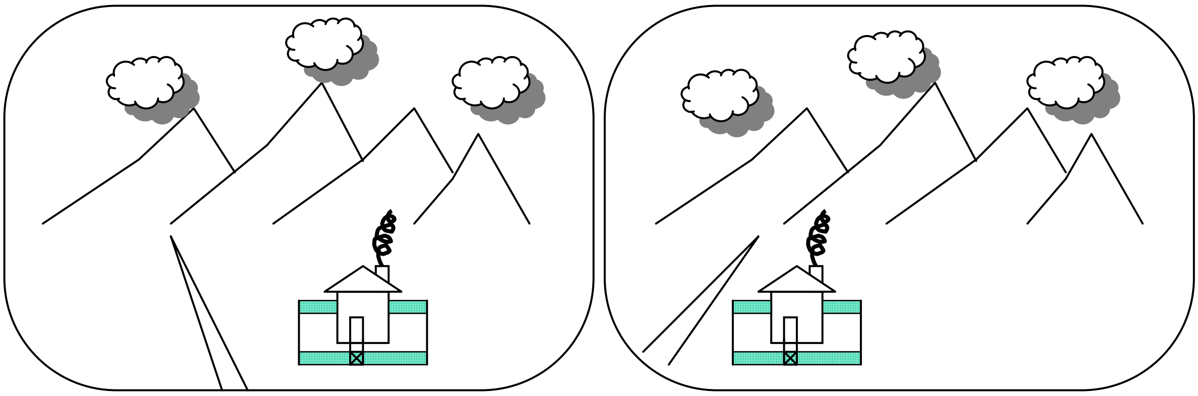

Figure 9-4 Parallax in action: Relative position change of objects from a moving train.

Imagine you are looking out of the window of a moving train as it traverses a plain with mountains visible in the distance. Visible in the right foreground is a house on the plain. You close your eyes for a few seconds and then open them again. The view is similar, but the farm appears to have changed from right to left of the tallest mountain in the distance, because of the movement of the train that you are travelling on. This is a change in parallax. Objects closer to the observer move through the observer's field of view faster than more distant objects. This same effect occurs with aerial photographs, where the tops of buildings, trees or hills appear to move further across successive photos than their bases do. This difference in apparent movement is extremely useful, as it is directly correlated with the height difference between top and bottom. The higher an object is, the closer its top is to the aeroplane and the further it will appear to move, relative to its base, on successive photos.

The outcome of this effect is that by determining the difference in distance between the top and bottom of a feature, we can calculate how high the feature is. We need some other information to do this accurately, notably how high the camera was when the photos were taken, how high the base of a feature is above sea level and the focal length of the camera, but this is all relatively easy to collect, from field survey, mapping and/or the photo title strip (viz. Figure 9-6).

Figure 9-5 Example of a Black and White (panchromatic) aerial photo with title strip information.

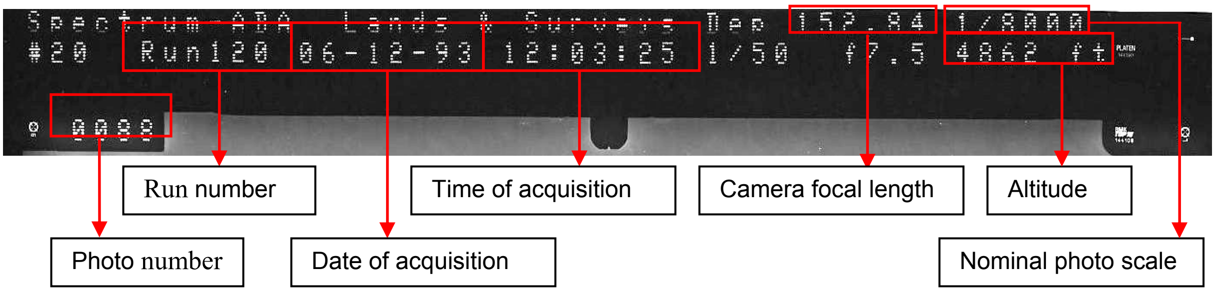

Figure 9-6 Zoom of the titling strip for the photo shown in Figure 9-5.

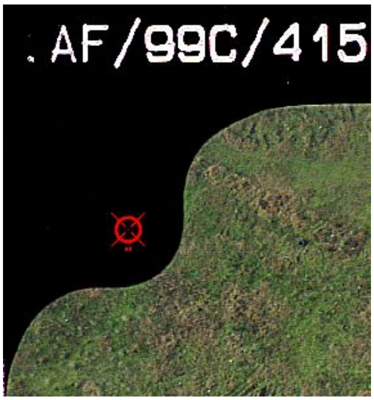

Figure 9-7 Example of a fiducial mark at the corner of an aerial photograph. There can be four or eight of these marks at the corners and in the centre of edges of photos, dependent on the camera make and model.

9.2.3 Scale in aerial photography

Perspective distortion means that photographs cannot have an accurate scale across an entire photo. The scale of an imaged object is dependent on how far away it is from the camera. Nevertheless, scale is widely quoted for aerial photography, but a user should be aware that a quoted scale is a nominal value – it gives an indication of the general level of detail and coverage of a photo, but should not be treated as a measurement ratio, as it can be with a map scale. Measurements derived using photo scales are likely to contain significant errors. If a scale for the photography is not given on the photo, or provided by the data supplier, a nominal scale is easy to calculate. There are a number of methods that can be used, but the most commonly used are:

- Direct comparison of the ground distance between objects visible in the imagery, usually calculated using a topographic map
- Calculation based on camera focal length and aeroplane elevation above the land surface

The first approach uses information from the airphoto and relates it either to a topographic map or to known real-world distances. The technique uses the following equation:

Equation 1:

$$S = \frac{Dp}{Drw}$$

Where: S = Scale of airphoto (nominal only)

Dp = Distance between objects on the photo

Drw = Distance between objects in the real world

When using the equation above, care must be taken to ensure that the numbers used for Dp and Drw are in the same units. As an example of how this equation might be used, imagine a photo where two path junctions are visible and also shown on a topographic map, the distance between the path junctions on the photo being 17.5 cm. The distance between the path junctions, as measured on our 1:50,000 map, is 2.4 cm. The equation then becomes:

$$S = \frac{17.5cm}{120{,}000cm} = 6{,}857$$

So our air photo has a nominal scale of 1:7,000, assuming the path junctions are at about the same topographic height (this could be checked from the map).

The second approach will give a nominal scale derived from three pieces of information, using equation 2:

Equation 2:

$$S = \frac{f}{H - h}$$

Where: S = Scale of the airphoto (nominal only)

f = Focal length of the camera

H = Aeroplane altitude, usually above Mean Sea Level (MSL)

h = Average altitude of terrain, above MSL

The aeroplane altitude and focal length of the camera are usually shown on the title strip of airphotos taken with a mapping camera. Because of the nature of these cameras, the focal length is not generally adjustable (as it is with most personal cameras), because the target is usually far enough away for the focal length to be set to infinity. The altitude should be recorded directly from an altimeter – if it is hand written, it may well be wrong! The average altitude of the terrain is sometimes difficult to decide, but an educated guess (based on a map or local knowledge) will allow photography of the most remote areas to be given at least a nominal scale. A great deal of care must be taken with units when using this equation as there are typically many conversions necessary. An example of this would be if (i) the focal length of a camera was given as 152.4 mm (a typical mapping camera, used for professional surveys, has a focal length close to this), (ii) the aircraft altitude was 4000 feet (4000 x 3.25 = 1231 m), and (iii) the average surface elevation was 150 m. Below is equation 2 with all values converted to metres:

$$Scale = \frac{0.1524}{1231 - 150} = \frac{0.1524}{1081} = 7093$$

So, the nominal scale of this photograph is also 1:7,000. In fact, it is entirely possible that the two values would come from the same photograph, as the scales are both nominal. There is every possibility that the differences are due to rounding error, unaccounted-for topography, error in the altitude reading, or a host of other error sources. Plate 15 shows the effect of topographic changes on an aerial photographic survey. If the aeroplane flies at a nominal altitude above sea level and the surface elevation changes, the scale of the photography will change significantly, too.

9.2.4 Further manual photogrammetry

Air photo surveys typically acquire photos with at least 60% overlap in the flying direction (Plate 16: upper) and at least 20% overlap (sidelap) between flying runs (Plate 16: lower). This ensures total stereo-coverage of an area so that photogrammetry can be undertaken between any points within the area. An entire area survey is usually termed a photo 'block'.

Many other equations exist for undertaking a host of calculations from aerial photographs, including parallax equations to allow an interpreter to calculate the heights of objects. Hardware to speed up this kind of calculation also exists, such as a parallax bar, which can reduce the complexity of each calculation between objects on a pair of air photos and is usually included with mirror stereoscopes. Stereoscopes allow an interpreter to view airphotos in such a way that the interpreter's brain is fooled into thinking a real surface is being seen in three dimensions. A stereopair of airphotos is viewed through a set of angled mirrors: when positioned correctly, the operator will perceive a three-dimensional view of the photographic subject. In this way, complex volumetric relationships can be assessed and evaluated. An example of a lab-spec mirror stereoscope is shown in Figure 9-8.

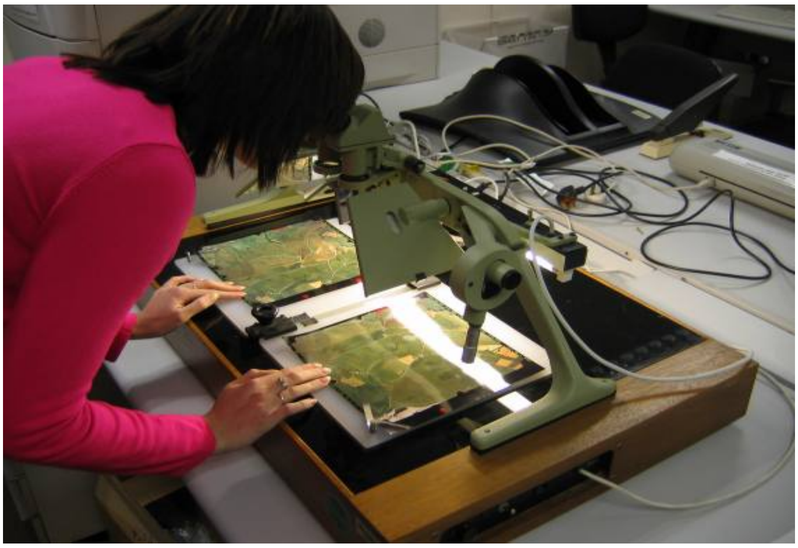

Figure 9-8 Air photo interpretation with a mirror stereoscope mounted on a light table (useful for viewing transparent air photos, diapositives, as well as traditional photographic prints).

Air photos can also be extremely useful in the field, especially if their 3-D viewpoint can be exploited. There are a number of ways this can be done, perhaps the most straightforward being the use of a pocket stereoscope (Figure 9-9). This equipment is small enough to carry into the field and will allow basic 3-D viewing of aerial photo stereo-pairs. Other approaches are use of virtual 3-D software such as ArcScene in ArcGIS and VGIS in Erdas Imagine, or the viewing of anaglyph imagery (two slightly different perspective images overlain in two different colours) in an appropriate software package on a laptop PC with red-blue glasses (Plate 17).

Figure 9-9 Pocket stereoscopes and aerial photo stereo-pair, one stereoscope set up for use, the other folded for transport. Photos are 9" per side for scale.

More complex photogrammetric machinery attempts to model the camera system used to acquire the photos, in an attempt to remove the distortions from the photography. The two main categories of these are analogue and analytical stereoplotters. *Analogue stereoplotters* are entirely mechanical systems for modelling the location, orientation and distortion of the camera system. Using a number of known survey points visible in the imagery, along with detailed information on the camera system, these devices allow an operator to construct an accurate, planimetrically correct map, directly from aerial photos.

Analytical stereo-plotters automate some of the above tasks, speeding up repetitive processes. Analytical stereoplotters are still sold today and are widely used for creating and updating maps. Both systems require an experienced operator and detailed photogrammetric knowledge to achieve good results - they are also very expensive to buy when new. Fortunately, software is now available that will allow photogrammetric processes to be undertaken without specialised equipment.

9.3 Geocorrection

Map co-ordinates are central to what makes a GIS different from a conventional information system – the spatial element to the data. A GIS allows us to spatially integrate data and interrogate disparate datasets based on the location of the features of interest. To undertake this with remotely sensed imagery, an image needs to be spatially located. Geocorrection is the process whereby imagery is assigned spatial map co-ordinates, transforming it into a valuable information source. Remotely sensed imagery and aerial photography are traditionally supplied without a digital co-ordinate system, so they need to be geo-corrected before they can be fully used in a GIS. This is changing, however, and some data, such as Landsat ETM, are now supplied as standard in a geo-corrected form.

The first stage in Geo-correcting an image is to select an appropriate map reference system for the area of interest. Correction of all data sources to a single map projection will allow accurate and easy integration, so choosing the correct system is critical to success. Refer to Chapter 2 for a more in-depth discussion of map projections and reference systems. The process of geo-correcting an image can be thought of as two-stage process, with transformation followed by resampling.

9.3.1 Transformation

To see how this process works, an understanding of file co-ordinates is necessary.

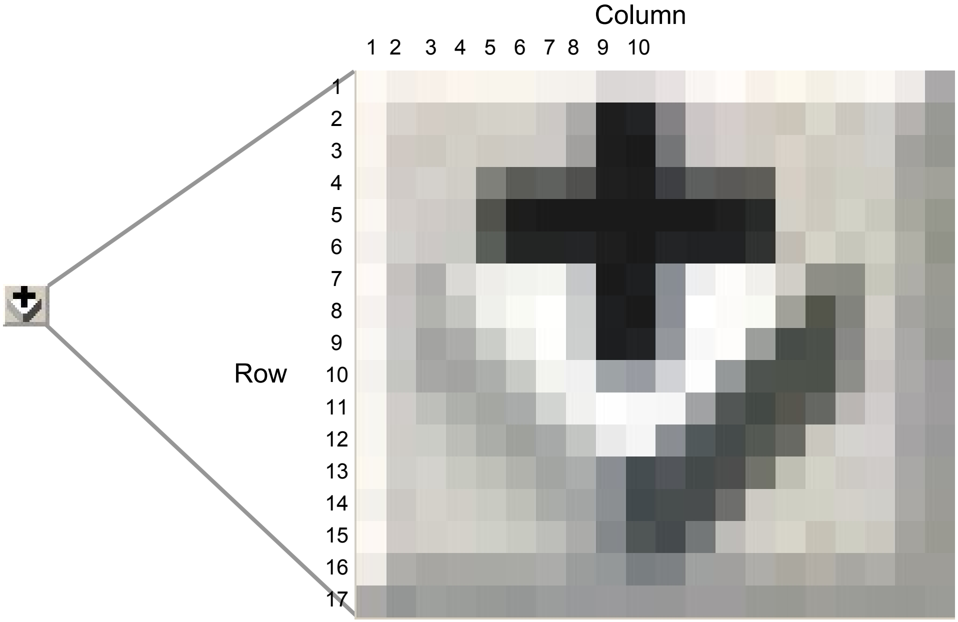

Figure 9-10 File co-ordinates of an image, usually given as column and row numbers. Note that pixel ‘1,1’ is at the top left, rather than bottom left.

File co-ordinates (as in Figure 9-10) reference each pixel in an image individually. If the image is a remotely sensed image of the Earth's surface, the pixel also represents an area on the ground for which map co-ordinates can be found. The process of transformation determines the mathematical relationship between a pixel's file co-ordinates and the map co-ordinates of the area that pixel represents. This is represented graphically in Figure 9-11.

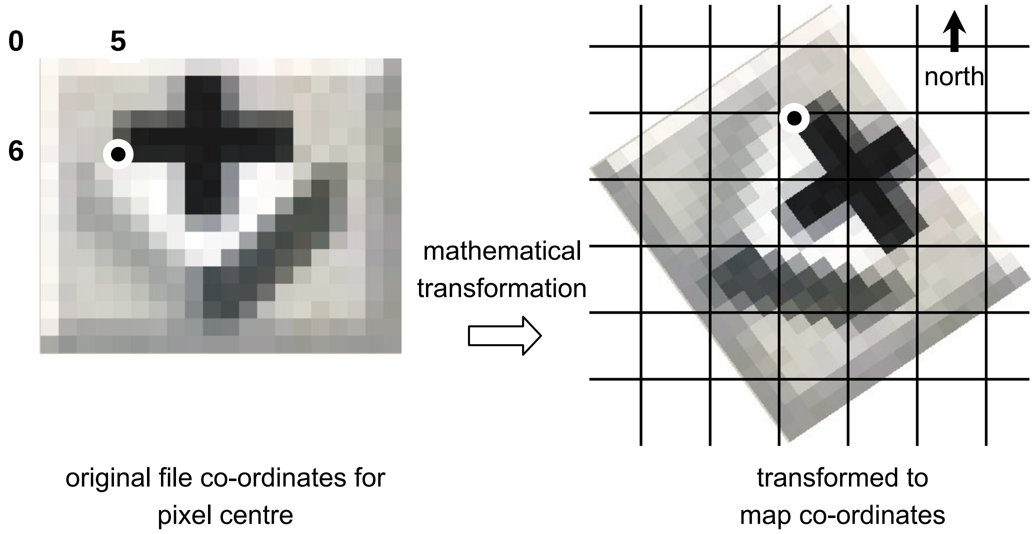

Figure 9-11 File co-ordinates transformed mathematically to map co-ordinates.

The choice of which mathematical transformation method to use is critical to successfully geo-correcting imagery. A wide range is available, but a comprehensive description is outside the scope of this chapter. For further information on the range available, please refer to the background reading suggested elsewhere in this manual and the relevant software's tutorials and manual. The two most commonly used transformations are affine and polynomial.

Affine transformations apply simple scale, rotation or translation corrections. They require no ground control, but require the user to know what the shift required is in terms of degrees rotation, direction and distance to translate, and difference in scale. Each package which implements affine transformations does so differently, so the user will need to consult the help files of the particular package used.

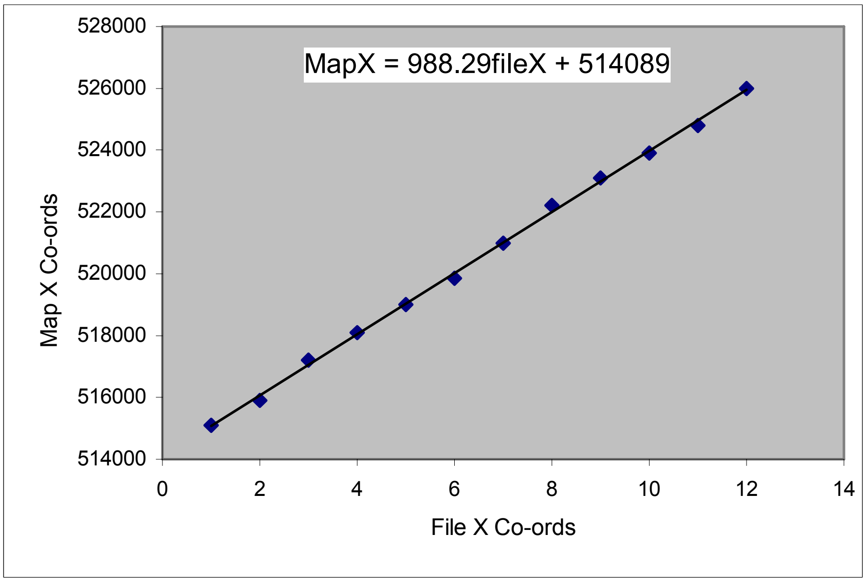

Figure 9-12 Map X co-ordinates (Easting) plotted against file X-axis co-ordinates with a polynomial best-fit line. Unknown co-ordinates can be estimated from the line position.

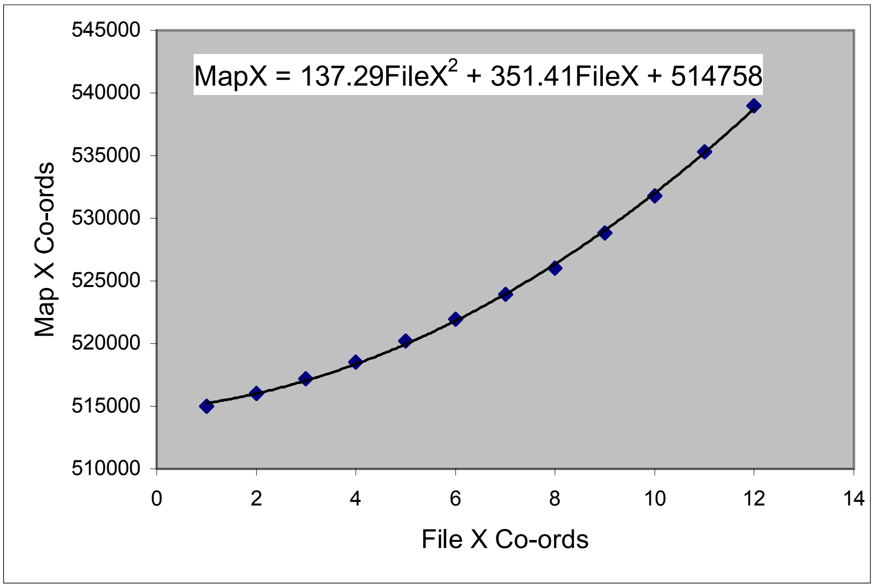

Figure 9-13 Map X co-ordinates (Easting) plotted against file X-axis co-ordinates with a second-order polynomial best-fit line.

Polynomial transformations are the most widely used method of mathematically transforming file co-ordinates to map co-ordinates. A polynomial transformation allows constants and variables to be applied to the original file co-ordinates, to obtain a logical desired result: the map co-ordinates. Conceptually, they can be thought of as a 'best fit' line through a series of points plotted on a graph. The defined constants and variables mathematically define this best-fit line, as illustrated in Figure 9-12 and Figure 9-13. Two equations are defined for each transformation, one for the X, Easting or Longitude co-ordinate and a second for the Y, Northing or Latitude co-ordinate. These equations are based on user-entered data and are used to estimate map co-ordinates for pixels not pre-defined by the user. Polynomial equations are defined by their 'order'. This is the number of variables allowable in the equation. The higher the order, the more complex the equation becomes (see Figure 9-14, Figure 9-15 and Figure 9-16). A higher order polynomial equation can remove more complex distortions from imagery than a lower order polynomial. Unfortunately, the higher the order, the more ground control is needed, and the greater the error as you move away from these points. Users are advised to use the lowest polynomial order that yields acceptable error levels to avoid these issues where possible. Which order to use depends on how the image needs to be distorted and in how many directions simultaneously.

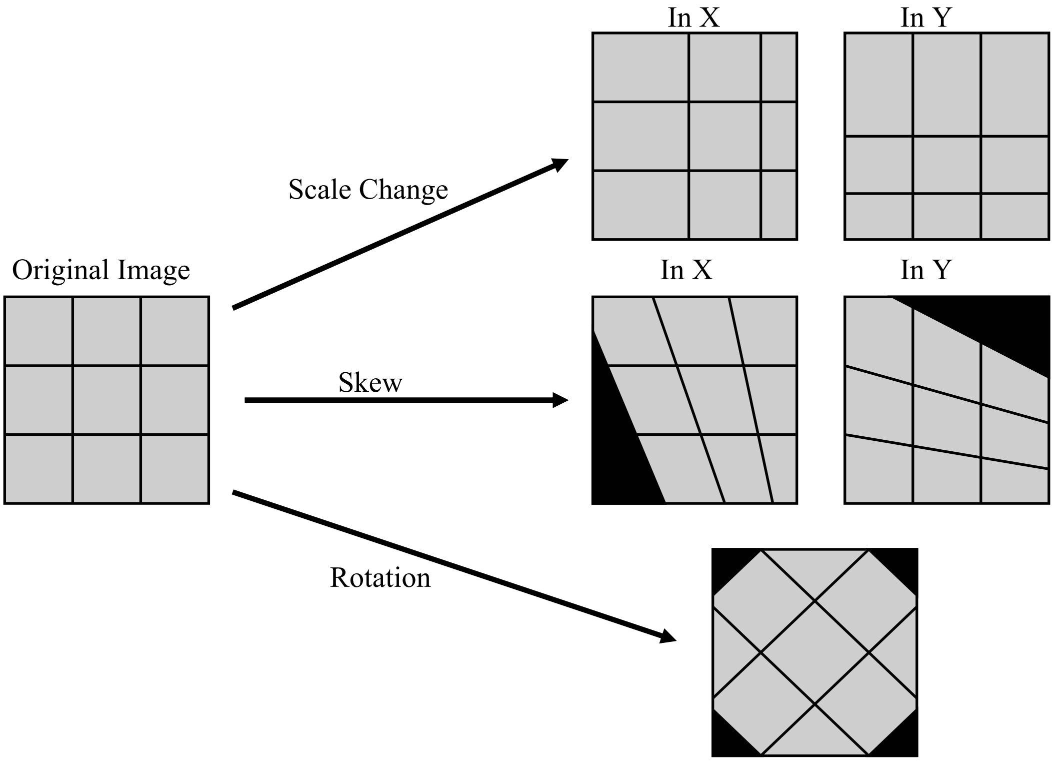

Figure 9-14 Distortions correctable with polynomial order 1 transformations.

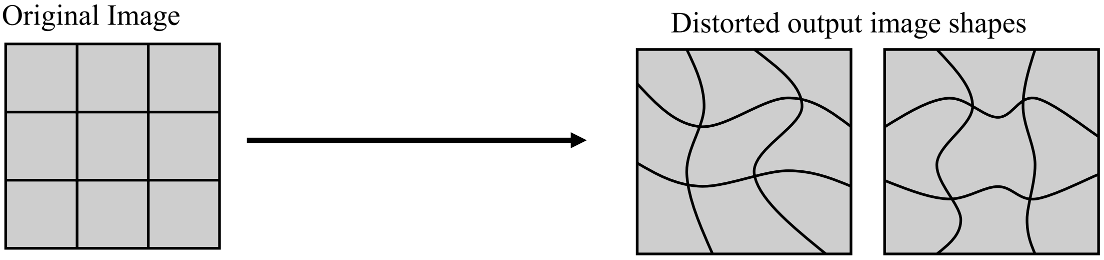

Figure 9-15 Distortions correctable with polynomial order 2 transformations.

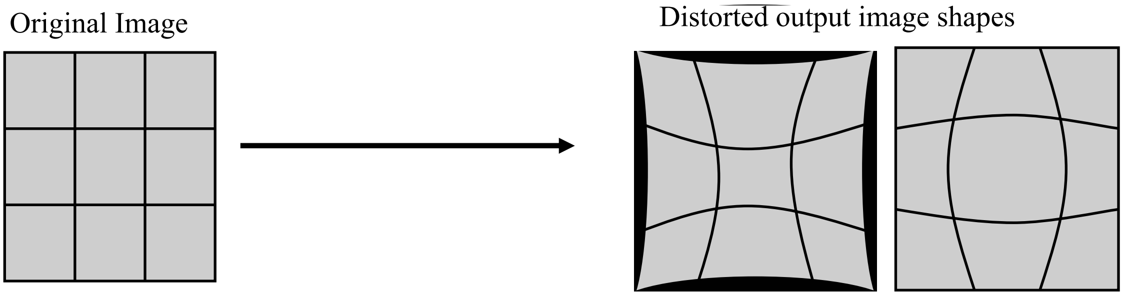

Figure 9-16 Distortions correctable with polynomial order 3 transformations.

9.3.2 Error measurement

The standard error measure in geocorrection is Root Mean Square Error (RMS error or RMSE). RMSE takes the difference between observed value and estimated value for each user-defined point, squares it, finds the mean squared value, and then finds the root of that value.

The equation for Root Mean Square Error is:

$$\text{RMS value} = \surd\, \Sigma(X_s - X_c)^2 + (Y_s - Y_c)^2$$

(where Xs and Ys are the user-derived source co-ordinates and Xc and Yc are the estimated 'best-fit' co-ordinates (see Figure 9-17))

RMSE is used in geocorrection because it gives a more sensitive and accurate measure of actual error in the result than simpler methods tend to. A modified calculation can also estimate individual Ground Control Point (GCP) error allowing the editing of points to improve the correction and reduce the error. RMSE is also widely used in scientific circles and, when correctly used, accurately describes the level of error we can expect to find in the corrected image.

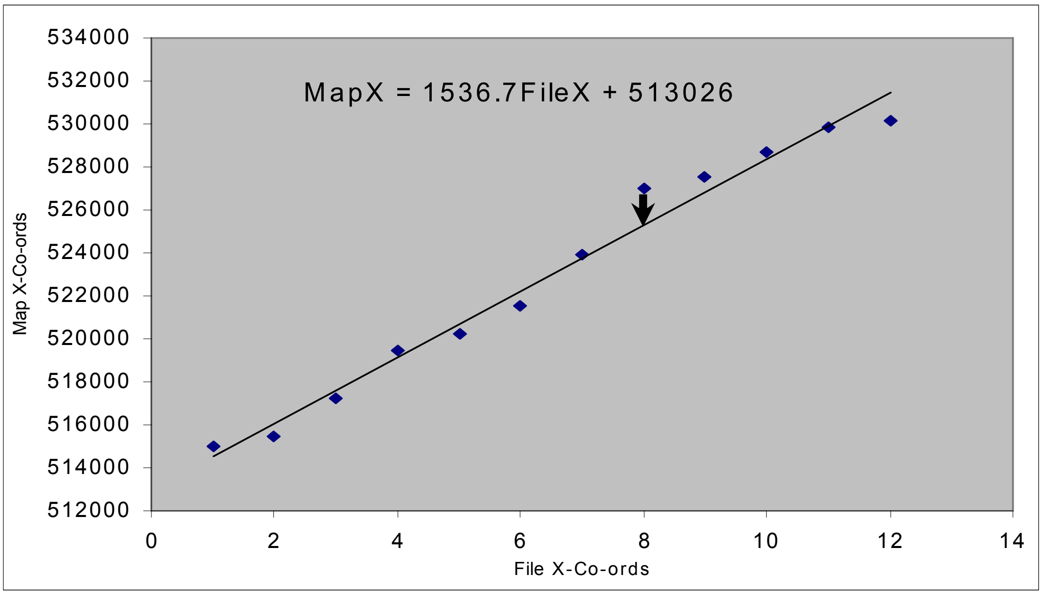

Figure 9-17 Arrow indicates the calculated error in the X-dimension. Total error is calculated by combining with Y-dimension error calculated in the same way.

Overall RMS error is calculated based on all the data entered by a user regarding the location of an image on the Earth's surface and the transformation method selected. Deciding what is an acceptable level of error should be based on an understanding of the process alongside a knowledge of the project for which the data is destined. RMS error values are defined in terms of *input image pixel size.*

The amount of error is related to the nominal pixel size of the image. For example, when correcting Landsat TM, the general rule is that the total RMS should be <1, but as a pixel is 30 m x 30 m, the actual error is up to 30 m. When correcting air-photos, the nominal pixel size may only be 25 cm. Even with an RMS of 7, the total error is still sub 2 m. Deciding on an acceptable error level therefore depends very much on the imagery, the quality of the data used to correct it and its anticipated use.

However, RMS error calculations are not foolproof: they are reliant on a good distribution of GCPs within an image. This may well be difficult, if not impossible, to achieve in some areas of the world. Poor distribution of GCPs can lead to huge errors that only become obvious when integration with other datasets is undertaken (Figure 9-18). Selection of appropriate and well distributed GCPs and choosing the correct transformation are therefore critical to limiting error in the output image.

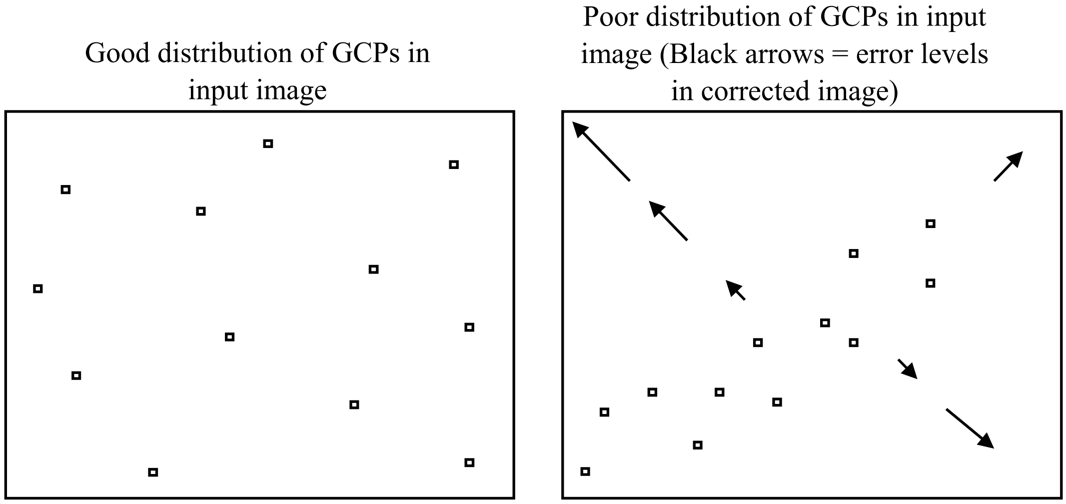

Figure 9-18 Effect on error in geocorrected image of GCP distribution. RMSE levels are the same for both images.

9.3.3 Resampling

The second stage of geocorrection is resampling, a process that creates a new image from the original, based on the information that you have entered. This stage is not essential for using the image, but computers deal very well with regular, orderly grids of square pixels, arranged to match a 'graph/origin' structure. Computers struggle to deal with data arranged in other ways, particularly with distorted grids and irregularly shaped pixels. Resampling allows an image that is distorted (relative to a map grid) to be stored in a regular, orderly grid of pixels, as illustrated in Figure 9-19.

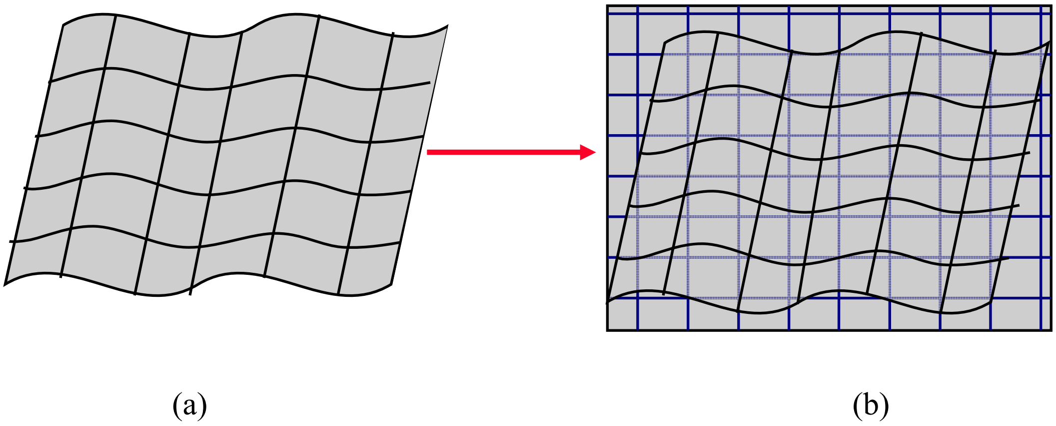

Figure 9-19 Creation of undistorted image grid via mathematical transformation and resampling of original image. (a) Distorted original image pixels. (b) New, regular, ordered image pixels.

The problem that resampling solves is that our transformed pixel centre co-ordinates might not fall exactly in the centre of a pixel in our new image. How do we therefore choose a value to represent our new pixel? Resampling provides the answer. There are three widely used resampling strategies available in most geocorrection-capable software: (i) Nearest-Neighbour, (ii) Bilinear Interpolation and (iii) Cubic Convolution. These three approaches for deciding what digital number to use for our new pixel center all have advantages and disadvantages. How these approaches work and issues with using them are discussed below.

The *nearest neighbour* approach uses the closest old pixel centre value as the new pixel centre value. The advantage with this is that the pixel values remain unchanged, which can be critical if further processing, such as image classification, is to be undertaken. A further advantage is that the decision rule used is simple to understand. The primary disadvantage is that a given old pixel centre might be closest to more than one new pixel, so the values are used twice, whilst another centre may not be used at all. This can result in ‘steps’ in images, irregular breaks in image features that are clearly artifacts of processing. The artifacts this process can produce can detract significantly from an image used for manual interpretation.

Bilinear interpolation calculates the new pixel value based on a weighted average of the four closest old pixel centers. The advantage with this technique is that processing artifacts, particularly ‘steps’, do not occur, giving an improved visual appearance. A second advantage is that the technique is computationally fast compared to cubic convolution. The main disadvantage is the original pixel values are lost, so further automatic image analysis is unsound.

The *cubic convolution* approach calculates a new pixel value, using the closest sixteen surrounding old pixel values. The main advantage with this approach is that it produces the most visually acceptable image, so is optimal for further manual processing. The main disadvantage, as with bilinear interpolation, is that the original pixel values are lost, making further automatic image processing unsound. The technique is also very complex to understand and computationally intensive.

Resampling is a fundamental process required to geocorrect remotely sensed imagery, however, the approach used can limit the utility of the corrected imagery. Even if imagery is supplied in a geocorrected form, knowledge of the processing applied to achieve this can be useful, as geo-corrected imagery will have to have been resampled at some stage. Care should be taken to ensure a user fully appreciates the pre-processing that an image has undergone, so that errors are not made.

9.3.4 Choosing Ground Control Points

Typically, to geocorrect imagery, a minimum number of points must be identified in the image, for which map co-ordinates are known and which are identifiable in the image (ground control points, GCPs). The number necessary will depend on many factors, but the main criteria are the transformation technique used, the type of imagery to be corrected and the accuracy required. Choosing optimal ground control points will aid in ensuring high accuracy and may also potentially reduce the number of points needed to correct an image. Points which are as precisely located as possible should be chosen. Road/path junctions should be selected where the angular relationship between the features is as close to 90° as possible. Figure 9-20 illustrates the issues with angular relationships when identifying points.

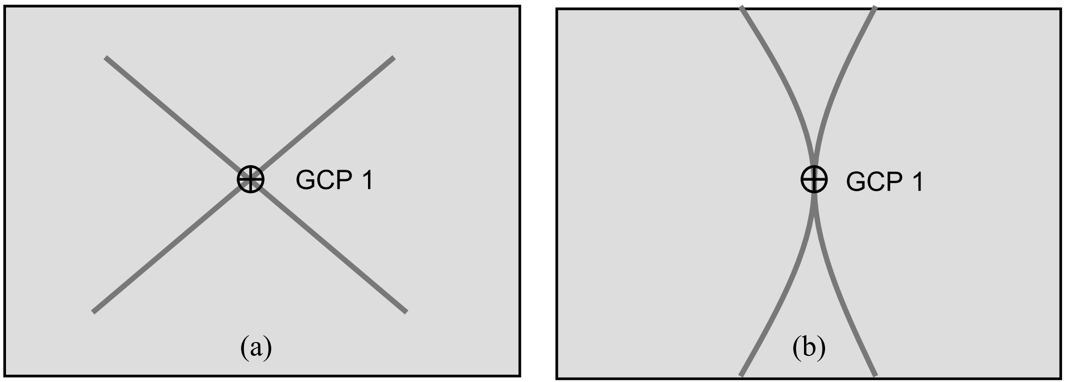

Figure 9-20 (a) Optimal and (b) sub-optimal location of a ground control point. (b) is less precise because the region in which the point could be correct is larger.

When using the polynomial transformation, there are theoretical limits on the number of points required to correct an image. With order one, at least three points are required. With order two, six points are needed and with order three, ten points are necessary. Many more should be used, however, to ensure reasonably accurate error reporting and redundancy in points. At least *four times* the minimum number should be defined before any faith should be put in the accuracy of the result. These should be well distributed throughout the image to ensure any error is minimised. When using ground control derived from GPS data, care must be taken to account for and assess the error inherent in the GPS derived co-ordinates. See the chapter on GPS in this manual for further information on GPS error.

9.3.5 Practicalities of geocorrecting images

The procedure for geocorrecting images is: first define which image you want to correct. This can be problematic, especially if the image is of an area of high relief, as geocorrection may give very poor results and orthocorrection becomes the only accurate option. In areas of low relief or where relatively large errors are acceptable, geocorrection will yield relatively reasonable results in most cases. The choice of which correction process to use is with the user and the choice will depend somewhat on their experience and the use to which the image will be put. Even the best quality satellite data may contain uncorrected distortions so a user must always be wary.

Second, decide what data to use for ground control: typically this will be a map, but could also consist of GPS co-ordinates and a description file, or other imagery, if it is already geocorrected. Digital map data can be used if available, although this may be difficult to find and purchase. Care should be taken that the data used for ground control is at an appropriate scale for geocorrecting your image. In many cases, small scale mapping will be all that is available. Features on this kind of map are usually highly generalized, so care is needed in selecting features that can be accurately located, and the error levels of using this kind of data should be kept in mind. Remember, 90 m on the ground is less than 1 mm on a 1:100,000 map, but represents three Landsat pixels, a significant error margin. Geocorrection is usually possible, regardless of the data quality, but the location error must be appreciated when using the output if serious mistakes are to be avoided.

The third step is to choose a transformation type. As already mentioned, numerous types exist, but the most appropriate will depend on the individual circumstances. Choosing a polynomial transformation is a good place to start unless an alternative, offered in the software, has been researched in some depth.

The fourth stage is the most time consuming, identifying and defining the ground control points. These need to be individually identified in the image, then their co-ordinates entered, from whatever source they have been derived. Control can usually be entered from another image, a digital map, from a text file (e.g., GPS co-ords) or from the keyboard, among others. Usually, entry from the keyboard is most prone to error. A single figure entered incorrectly can place your point thousands of kilometers out. Real care needs to be taken that every value entered is correct. Real time error reporting, shown in Plate 18 as a column for individual point RMS Error, is a great help in identifying the problem.

Point filtering is the final step in the process. The points entered should be reviewed to identify any serious errors. This could be due to operator error, but could also be due to map generalization. Experience in this area is essential, so the new user is advised to refer to a more experienced colleague, but in the absence of this advice, common sense and careful contemplation can suffice. An example of point errors can be drawn from a personal experience with Spanish 1:25,000 mapping when correcting a Landsat scene. A road over a stream marked on the study area map was identified as a likely GCP. Having placed the point on the image and entered the co-ordinates for the crossing, the error level was >3, which was unacceptable for the intended use. When the map was re-examined, it was found that the stream marked on the map did not actually cross the road where the bridge was marked. This led to a re-examination of other areas on the map. It turned out that the map had been mis-printed and consequently the colour layers were not precisely registered. All inter-relationships between roads and waterways were mis-located. When points using this kind of relationship were removed, the error fell to <0.7.

When maps are produced, more care is usually taken to place anthropogenic features accurately than natural ones. River junctions should therefore be treated with caution, whereas major road junctions can usually be used with confidence. The final, but possibly most important point about GCP filtering is to resist the temptation to automatically remove the points with the highest error. This can very easily lead to poor point distribution, which can produce a less accurate output than a high error level with a good distribution. This is especially true of isolated points towards the edge of the image, these can be essential to maintaining the overall positional accuracy of an image.

9.4 Digital photogrammetry

Digital photogrammetry uses the power of computers, namely rapidly repeating complex calculations, to reduce the difficulty of producing ortho-images and surface elevation models. By using software for this kind of processing, the requirement for expensive and complex machinery is substantially reduced. A little explanation of the terms ortho-image and surface elevation models is, perhaps, appropriate here. Ortho-images are images that are planimetrically correct and so can be used directly as maps. All distortions have been removed, so scale is constant, angular relationships are correct and areas can be measured directly. Surface elevation models are representations of (usually, although not always) the Earth’s surface. They can consist of line data (contours), regularly spaced spot heights

(Digital Elevation Model or DEM) or irregularly spaced linked spot heights (Triangulated Irregular Network or TIN).

Production of ortho-photos and surface models from aerial photography involves modelling the location, orientation and optics of the camera in relation to a map co-ordinate system. Knowing precisely how the camera is located in three dimensions and how light is transmitted through the optics onto the photographic film, allows the software to retrace the light that formed the photographic image. From multiple viewpoints and from a few ground control points, the Earth's surface relative to the camera (using parallax – see last section) can be determined. Using this information, the full path of light, which is reflected from the Earth's surface, passes through the camera system and is recorded by the photographic film, can be modelled. Retracing this path allows the software to correct the location of each point in the image, resulting in a planimetrically correct image (Figure 9-21 and Figure 9-22).

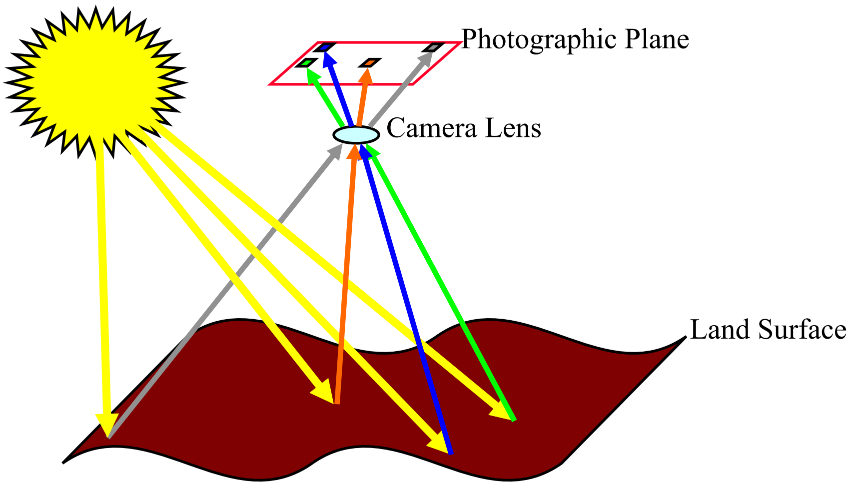

Figure 9-21 Conceptual diagram of an aerial photographic system. Land surface is sun illuminated, the reflected light passes through the lens, which focuses an image of the surface onto a 2-D photographic plane.

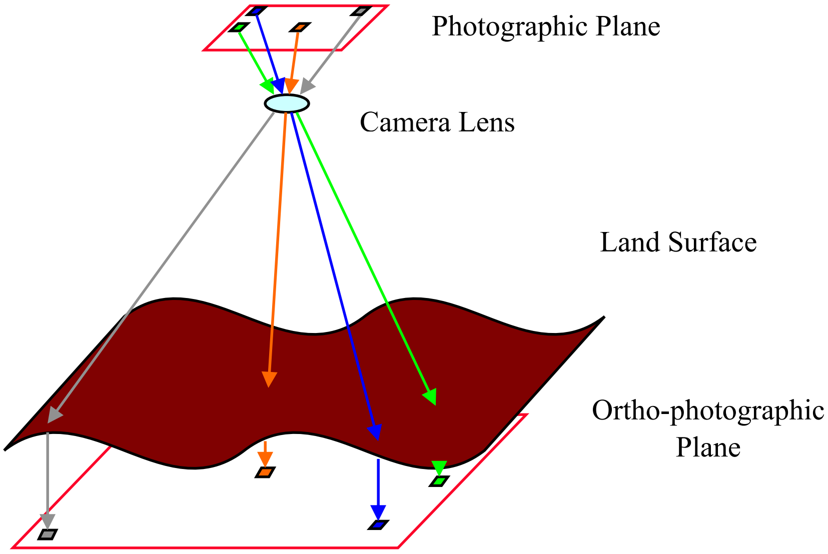

Figure 9-22 Ray-tracing approach to ortho-correction of aerial photographic distortion. The image is re-projected, via a surface model onto a planimetrically correct ortho-photo plane. All distortion is removed to produce an 'image-map', usually called an ortho-photo.

The result of this modelling procedure is extremely useful, giving all the information of a topographic map, but also including information usually excluded from traditional mapping, such as vegetation density, vegetation type and geomorphological context. The surface model created can also be extremely useful, giving information on a much finer scale than most maps, so that, for example, floodplain features are visible and can be analysed. The level of detail is only limited by the original photography and information not visible on standard maps is usually easily found (Figure 9-23).

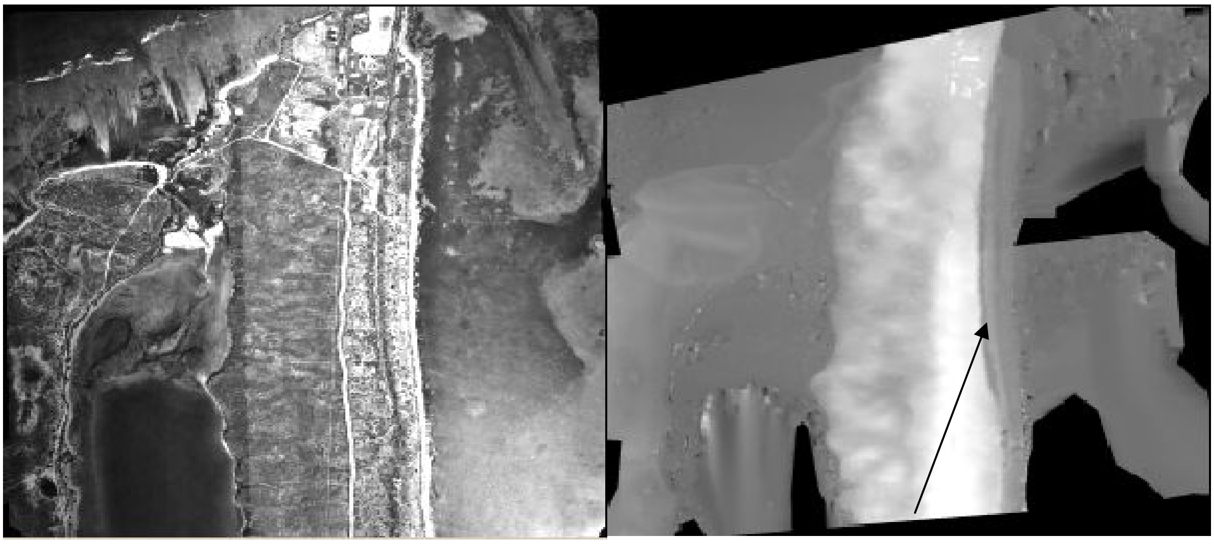

Figure 9-23 Orthophoto and DEM of the northern part of Grand Turk island, part of the Turks & Caicos island group. Note the beach forms on the eastern side of the island. The lighter band (arrowed) is a raised bench, probably caused by storm processes.

As all the output from digital photogrammetry is, by definition, digital, the image and the elevation model can be combined to produce pseudo-3D views and fly-throughs of the area of interest. Figure 9-24 shows a 3D view of Grand Turk, constructed from the images shown in Figure 9-23, illustrating the information present in the model and its usefulness in understanding the morphology of the Earth's surface.

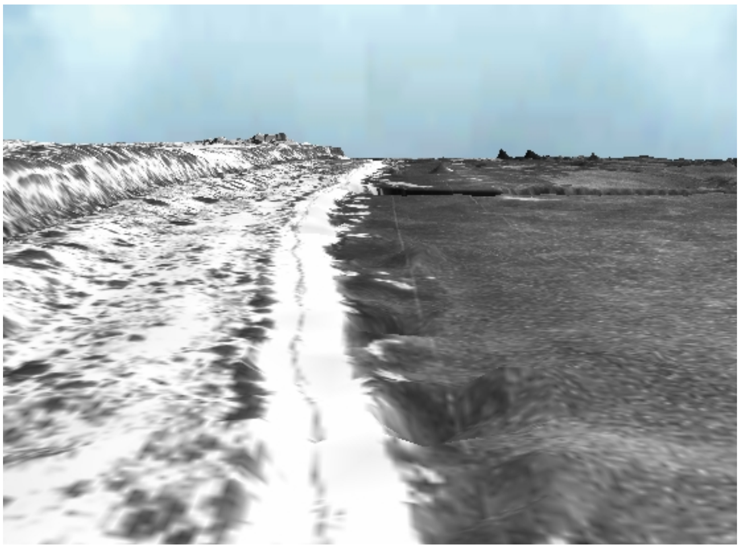

Figure 9-24 Pseudo-3D view of shoreline (arrowed in Figure 9-23) from the south.

This kind of information, which can be extracted from the photography without the need for field data of any kind (some ground information is necessary – see below) can be invaluable in locating features or areas of particular interest, finding accessible routes into remote areas, locating base camps in appropriate areas and for adding value and understanding to the fieldwork proposed. For conceptualising the terrain of an area, this approach is unsurpassed.

9.4.1 Undertaking Digital Photogrammetry

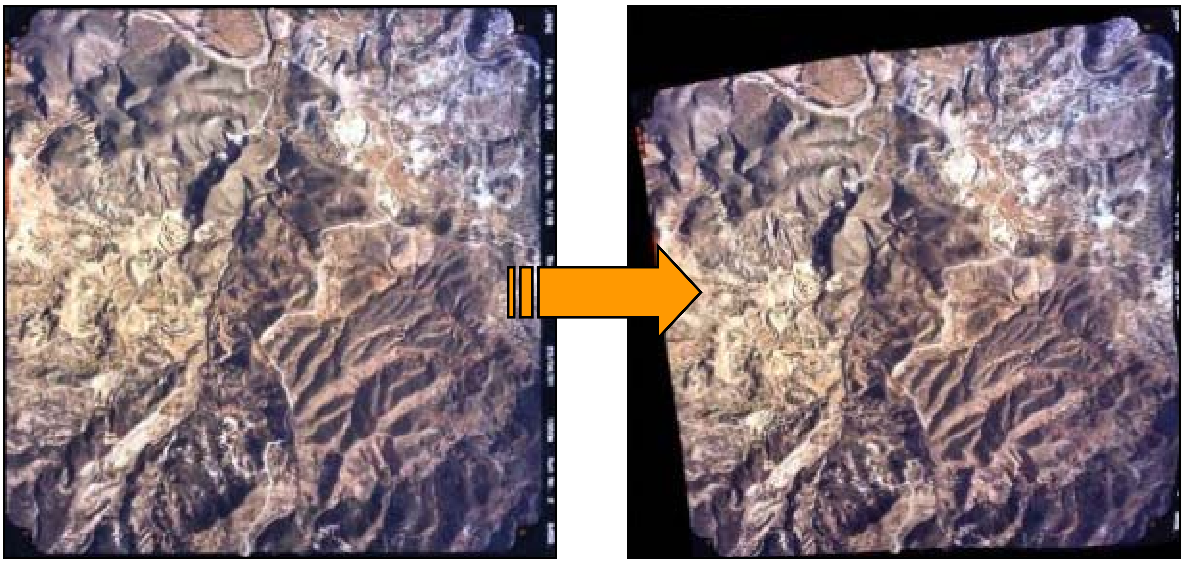

Figure 9-25 The image on the left is an uncorrected scanned air photo, the image on the right is an orthophoto produced from digital photogrammetry of the same area. Note the distortions in the shape of the photo caused by removal of camera distortions, topographic distortion and perspective distortions.

Digital photogrammetry software is quite specialised and complex to use. Detailed descriptions of the processing stages are outside the scope of this manual and you will need expert advice and assistance to undertake the process yourself. If you are aware of the necessary stages, however, it is possible for a committed novice to achieve a successful result using information from many sources, particularly software manuals and the internet.

The first step in the process is usually to get the photography into digital form and a number of issues need to be addressed at this point in the process.

Most aerial photography is analogue, i.e. the photo is recorded on film using chemical reactions. Analogue photographs must therefore be scanned before they can be used digitally. To achieve the best results, particularly in terms of subtle elevation changes, a photogrammetric-quality scanner should be used and film-based diapositives should be scanned rather than the traditional paper prints. If a high accuracy (usually <50 cm vertical error from 1:10,000 photos) is not essential, paper prints scanned in a reasonable quality desktop scanner are more than adequate. The photos should be scanned at an appropriate resolution for their intended use – generally 300-600 dpi is appropriate. The more detailed the result required, the higher the resolution should be, but care needs to be taken that the optical resolution of the scanner is not exceeded as this adds significantly to the data file size without increasing the actual detail visible in the image.

The next stage would be to import the scanned photos into the software package. This is typically done as a photo 'block'. Usually this block would consist of all the photos that are necessary to provide stereo-coverage of the area of interest. The photo block may consist of multiple photos per run and multiple runs to cover a wider area. All should be processed simultaneously to achieve the best results. In practise, this may not be feasible due to storage constraints, particularly if the images are scanned at a high resolution and in colour (a single colour photo, scanned at 600 dpi, requires around 80 MB of storage. Processing it will take around three times this space. To process 20 photos would therefore need at least 5 GB of storage).

At this point you would need to define the internal parameters of the camera used to acquire the photos. This information can be found on the Camera Calibration Certificate, which should be available from the same source as the photos. Most software has default camera specifications that can be used, but this will mean some camera-induced distortions may not be properly removed. At the very least, the information about the camera focal length should be correct. This information is generally included on the photography itself on the titling strip.

Once the camera has been defined, the next stage is to define the internal orientation. This means matching the camera model to the digital images and simply involves defining where on the digital photos the fiducial marks appear. These marks are precisely defined in the camera calibration certificate, so can be used to determine where any part of the image was located on the photographic plane when the photo was taken.

The next stage is to define the exterior orientation. In this process, the position of the camera in relation to the chosen map co-ordinate system is determined through the use of ground control points. Further points can be defined, called tie points, which are not linked to the map co-ordinates, but help to 'tie' a stereo-pair of photos together so that ground control on one can be used to correct both. Once a series of ground control points (GCPs) and tie points has been defined for a whole photo block, triangulation can be undertaken. As a guide, a minimum of one GCP and four tie points per photo is required, though clearly more points would be useful in reducing error. Triangulation is an attempt to reconcile the camera model, the ground control points and the point locations to an acceptable model. If errors in data entry have occurred, triangulation is likely to fail. Error reporting at this stage is a good indicator of the error present in the final datasets. Time spent reducing error levels at this point is usually well spent, in terms of improved accuracy in the DEM and ortho-photo output.

The penultimate stage is DEM extraction. The DEM for each stereo-pair in turn can be derived for the area of stereo coverage through reference to the triangulation information and pattern matching between the images. The pattern-matching algorithm attempts to match the two photos together, based on the nearest GCP or tie point to a particular pixel. When the algorithm finds a match, a parallax measurement is undertaken which calculates the topographic height of that point, based on the GCP data. This height is recorded and the algorithm begins to try to match the next point on the image. This continues across the whole area of stereo-coverage, building up a digital elevation model. Some software is capable of directly deriving a TIN via a similar method, but the process is fundamentally the same.

The last stage is *ortho-correction* of the photography. Having produced an elevation model, all the necessary components for the ray-tracing approach are defined, so the process re-projects the photo image onto an ortho-photo plane, removing all distortions and producing a planimetrically correct image. Ortho-correction provides accurate image-maps that can be used on the expedition in the same way as traditional topographic maps.

9.5 Practicalities of accurate mapping using airphotos

Aerial photos in remote regions can always be useful, even if they are only used to orient yourself relative to ground features. Sources of aerial photography are discussed in the Remote Sensing chapter and contact details for major suppliers are given in the Appendix. For many projects, air photos will not be available, or ground control will be insufficient to fully gain the benefits they can provide. For other projects, a reconnaissance trip would be needed but the financial limitations mean this is not an option. If air photos are available, they can provide invaluable information regardless of the level at which they are processed, so the effort to acquire them is rarely wasted.

Airphotos show land cover features that are missing from traditional topographic maps, and are usually much more detailed than available mapping. As long as the distortions inherent in photography are understood, photos can be used in a similar way to maps, helping find trails, locate potential campsites, navigate around areas and regions, etc, especially if used in conjunction with a stereoscope, whether pocket or lab-based. For more complicated tasks, undertaking ortho-correction and DEM extraction are likely to be necessary, especially if combining the data with other GIS data layers. This requires ground control, however, and getting this information can be difficult or even impossible. To fully correct a series of aerial photos, a significant number of features must be identified on the photos, each feature being assigned accurate co-ordinates and altitudes. On average, three ground control points per photo are required for an acceptable result, although these do not need to be evenly distributed over a sequence of photos. Some of the photos can be included with no ground control at all, as long as the other photos in the run have sufficient points. Collecting or finding accurate location information can be difficult: the best approach is to make a reconnaissance trip. Copies of the photos can be taken on the recce trip, features visible on the photos identified on the ground and GPS readings of the location taken to provide ground control. These ground control points can then be used to accurately ortho-correct the images.

This optimum of at least two trips to a study area may not, of course, be possible with all expeditions. With this scenario, obtaining ground control must be achieved through other means. GPS co-ordinates from other researchers who have visited the area could be used if there is sufficient information about where a reading was taken, although care must be taken to exclude ambiguous points. However, such data are always suspect, so caution must be used when using data not collected personally. Mapping, if it exists, can also be used, but much mapping around the world is very dated where it exists at all. Small scale maps (1:100,000, 1:500,000 etc) are insufficiently detailed to provide accurate control for most aerial photography, so 1:50,000 or larger scale maps are needed. Features derived from this type of mapping also need to be interpreted and their potential accuracy weighed. A good example of this is a recent project looking at a rainforest area in Guyana. The only ground control available was 1:50,000 maps, originally surveyed about fifteen years after the photography was taken. Very few features on these maps related to features on the photos, though some path junctions and river features were eventually used. One point used for ground control was a river confluence with a very distinctive sediment bar: this eventually turned out to be located around 300 m from its location on the map, probably because of channel migration or mapping error. As with all aspects of an expedition, advance planning can be critical to success, so plan your data production and use of aerial photography carefully.

9.6 Summary

The key with using all photography or imagery, whether aerial or satellite derived, is to only use the data in a sensible and meaningful way, understand the error levels and their potential impact, and plan a project in such a way that possible problems are avoided or mitigated where possible. *Geocorrection* can be a complicated process, full of pitfalls for the unknowledgeable and unwary. However, when undertaken by a knowledgeable user with care and attention it transforms a ‘pretty picture’ into really useful GIS information. Geocorrection is an essential stage in using remotely sensed data, and an understanding of the process helps the user to appreciate the limitations and advantages of the images they are using, even if they did not undertake the geocorrection themselves. *Photogrammetry* is a useful tool that everyone involved in understanding spatially distributed processes, objects or features should be aware of. For the explorer, understanding the terrain of the area to be explored is clearly critical, so photogrammetry can help them gain that understanding from photography prior to any visit to the region. For most projects, it is likely that having an understanding of photographic distortion would be sufficient to prevent any disastrous mistakes, but for more detailed survey work, more in-depth understanding and knowledge may well be necessary.