Title	Digital elevation modelling of individual monuments from aerial photographs
Author(s)	Redfern, Sam; Lyons, Gerard J.
Publication Date	1999
Publication Information	Redfern, S.; Lyons, G.; Redfern, RM; (1999) 'Digital elevation modelling of individual monuments from aerial photographs'. Archaeological Prospection, .
Publisher	John Wiley & Sons
Link to publisher's version	http://dx.doi.org/10.1002/(SICI)1099-0763(199912)6:4<211::AID-ARP125>3.0.CO;2-7
Item record	http://hdl.handle.net/10379/4073
DOI	http://dx.doi.org/10.1002/(SICI)1099-0763(199912)6:4<211::AID-ARP125>3.0.CO;2-7

SOFTCOPY DIGITAL ELEVATION MODELLING OF INDIVIDUAL MONUMENTS FROM STEREO PAIRS OF AERIAL PHOTOGRAPHS

S. REDFERN and G. LYONS

Information Technology Centre, NUI, Galway, Ireland.
Ph.: +353 91 524411 x 3327
Fax: +353 91 750501
E-mail: sam@it.ucg.ie

ABSTRACT

Digital elevation models (DEM) are crucial data products for a variety of geographic applications, and their generation directly from digitised stereo pairs of vertical aerial photographs has recently been accomplished. Despite the growing number of software packages providing softcopy topographic photogrammetry, there is still a need for practical approaches which do not require accurate lens calibration information or time-consuming ground control.

This paper presents an algorithm for generating small-area DEMs directly from digitised aerial photographs, with no additional information required other than the flying height of the aircraft and a small number of control points measured from maps.

KEYWORDS: DIGITAL ELEVATION MODELS, SOFTCOPY TOPOGRAPHIC PHOTOGRAMMETRY, REMOTE SITE SURVEY.

INTRODUCTION

Digital elevation models (DEMs) are valuable data products for archaeologists studying individual monuments, as well as the topographic contexts of wider archaeological landscapes. DEMs, when used in a geographical information system (GIS) environment, allow a variety of useful studies, involving for example viewshed or area of visual influence analyses, and calculations of slope, aspect (facing), and volume. The accuracy of a DEM is very important, because it controls not only the quality of the overall landform representation, but also the nature of all the secondary data derived from it (Kvamme 1990).

Professional photogrammetry equipment, capable of generating DEMs from stereo pairs of overlapping vertical aerial photographs, has existed for a number of years. However, its cost and technical complexity has effectively made it unavailable to archaeologists. The two DEM generation approaches commonly taken by archaeologists have been: interpolation from contour maps, and ground survey using equipment such as electronic distance measure (EDM) devices (Haigh 1993). The former of these techniques is both labour intensive and inaccurate, while the latter is extremely labour intensive, though very accurate. Photogrammetry, by comparison, allows good accuracy while being highly labour efficient.

Computer-assisted photogrammetry solutions have existed for a number of years: however, it is only within the last 3 years that reasonably priced solutions to PC-based softcopy photogrammetry have become commercially available. A photogrammetry approach is referred to as 'softcopy' when it is implemented entirely through software, without requiring additional specialist hardware. A number of software products now provide the required functionality, yet the archaeological community has not yet widely availed of them. This appears to be due to several factors, including:

- A lack of awareness of these products;
- The fact that accurate knowledge of the cameras and photographs is typically required, for the photogrammetric equations of interior orientation. Since this information is not recorded on the photographs themselves, it is often unavailable for photograph sets that were not commissioned by the archaeologist users;
- The fact that precise ground control is typically required, for the photogrammetric equations of exterior orientation, which hinders the potential benefits of using stereo aerial photograph sets for rapid wide area survey;
- The fact that the available software applications, though more accessible than traditional photogrammetry equipment, still presume a strong technical background in the areas of photogrammetry and computing.

This paper provides an introduction to the techniques typically used by softcopy photogrammetry systems, and presents a practical new approach, specifically developed for archaeological purposes. The overall accuracy of this approach is quantitatively assessed, and its contribution more generally to softcopy photogrammetry is suggested.

SOFTCOPY TOPOGRAPHIC PHOTOGRAMMETRY

Parallax

Fundamental to analytical photogrammetry is the principle of *parallax*, which is the apparent shift in the position of an image due to a change in viewing position. Objects or points in the landscape which are closer to the viewer will appear to move faster than objects that are further away. For the overlapping region between successive vertical aerial photographs, the amount that each point in the landscape is displaced from one viewpoint to the next can be used to compute the height of that point. The measurements made are of the parallax apparent parallel to the direction of flight of the plane, and can be made between any clearly distinguishable points in the photo overlap. For truly vertical photographs taken from the same flying height, points at the same height will theoretically have identical parallax values, and the parallax value of each point will be directly proportional to its height (Slama *et al.* 1980). In practice, geometric image distortions due to lens and print imperfections, as well as due to camera tilt, render this assumption invalid.

Digital Image Processing Fundamentals

Using a scanner, photographs are imported into a computer as bitmaps - grids of square pixels each of which is stored as a single number representing a discrete colour or greyscale intensity (figure 1). Digital image processing techniques such as softcopy photogrammetry are based on the numerical manipulation of these pixels - often millions of them - and the analysis of the spatial relationships between them.

Figure 1 Part of a digitised aerial photo with a section containing a ringfort enlarged to show its pixels.

The bulk of the computational effort performed by any digital photogrammetry system is the automatic matching of thousands or millions of corresponding pixels between successive images, in order to determine the parallax value at each pixel. The most common technique used is called *cross correlation*. This is a technique whereby a computer searches for a sub-image or pattern in a larger image. A simple way of implementing this is to take the sub-image and test it at each possible position in the larger image, to see the position at which it fits best. For example, if one were searching for bright squares with an area of 9 pixels on a dark background, the sub-image shown in figure 2 could be used as a template. To search for this sub-image, the template is moved across the large image, and at each position the absolute differences between the template pixels and the values of the pixels directly below them are summed. This will result in a coefficient of correlation at each position: the lower this number, the better the fit at that position.

0	0	0	0	0
0	255	255	255	0
0	255	255	255	0
0	255	255	255	0
0	0	0	0	0

Figure 2 A cross-correlation template for finding bright squares (pixel values 255) on a dark background (pixel values 0).

In order to determine the parallax value of a pixel in a photograph from a stereo pair, a cross-correlation template is established using the numerical values of the pixels in a given neighbourhood around the pixel under analysis. This template is then applied to the second photograph, in order to find the

best match. Cross-correlation is a time consuming and inexact process when the search pattern does not exist perfectly in the search image, which will certainly be the case when working with stereo aerial photographs. Therefore, the strategy used to determine the search region, that is the region in the second image that is to be searched, is crucial.

Published Softcopy Photogrammetry Solutions

A number of algorithms for softcopy photogrammetry have been discussed in the photogrammetry literature, though increasingly the trend has been not to publish specific implementation details, probably due to the commercial potential of these algorithms with the advent of the desktop PC and of GIS software.

The analytical steps carried out by most digital photogrammetry systems are (see for example Krzystek 1995; Dupéret 1996):

1. *Determination of exterior and interior orientation of image*. The camera's position in Cartesian space and its attitude (tilt, swing, azimuth) and inherent distortions are calculated. This requires knowledge of calibrated focal length, radial lens distortion, and principal points, as well as a number of precise ground control points (see Wong 1980);
2. *Cross-Correlation* or *Disparity Matching*. This involves the estimation of the x-parallax values of some subset of the image pixels, typically using the cross-correlation coefficient method;
3. *Blunder/noise detection and removal.* This step aims to detect all erroneous parallax values, and is typically user-driven. It is regarded as being one of the primary issues still to be solved for fully automated systems;
4. *Calculation of heights from x-parallax values*. The application of the image models calculated in the first step allows parallax values to be linearly translated to height values;
5. *Interpolation or re-sampling*. The final step is normally to modify the output values in order to generate a DEM of user-specified grid size. This may involve interpolation to fill in missing values, and/or averaging to reduce x/y resolution.

Search Region Reduction

The bulk of digital photogrammetry research has been concerned with the stereo matching stage: the main concerns have been speed and accuracy. Cross-correlation mis-matches are minimised by reducing the search region in the second image. The most common approach is to form an image pyramid (hierarchy of increasingly high resolution stereo images), which

allows the generation of progressively higher resolution DEMs. Each intermediate DEM is used to narrow the search region for cross-correlation in the next (see for example Ackermann & Krzystek 1995). This technique is used for two reasons: at low image resolutions, high frequency noise affects correlation coefficients less than at high resolutions; and, physically larger image features may assist the low-resolution searches.

While this "coarse-to-fine" image pyramid approach is the most commonly used technique for search region narrowing, others are not unknown. Search region reduction may also be implemented by applying a knowledge of the approximate overlap area and the average height of the terrain (Heipke 1996). Goshtasby *et al.* (1984) propose cross-correlation in two stages, primarily in order to speed up the matching process. The first stage determines an approximate match using a small correlation patch size: this is used to direct the correlation of the second stage, which uses a larger patch size.

Texture Improvement

Featureless local areas with very poor texture cause problems for cross-correlation based systems, which may produce prohibitively large amounts of mis-matches. A common solution is to apply contrast improvement. Figure 3 illustrates the very poor quality results obtained from low texture images without texture improvement.

Figure 4 illustrates the improvement achieved when the stereo images are modified using histogram equalisation contrast improvement. Histogram equalisation spreads out the greyscale pixel values of an image, in order to produce an equal frequency of each available value (i.e. equalise the image's histogram of greyscale frequencies), and therefore maximise the overall contrast.

Figure 3   DEM of a ringfort generated without improving the texture of the stereo images (left) one of the images of the ringfort; (right) DEM generated from a single correlation pass with a 9x9 window - this is almost entirely noise.

Figure 4   Texture improvement using histogram equalisation. (left) Histogram equalised ringfort image; (right) DEM generated from a single correlation pass with a 9x9 window, using histogram equalised images. Most of the pixels have now been correctly matched.

Obstruction, Blunder, and Noise Removal

The problem of obstruction, which occurs when the terrain is covered by obstacles such as trees, house and bushes, is regarded as being significant, particularly when working with urban scenes: this is one of the areas in digital photogrammetry still requiring further research (Ackermann & Krzystek 1995; Krzystek 1995; Ackermann 1996). The solution is normally to allow interactive editing by the user, which can be very time consuming.

Blunder and noise removal may also be carried out through averaging procedures (e.g. Al-Rajeh 1994; Kirsch 1997). Although some research papers mention the detection and removal of obstruction, noise, and blunders, explicit removal techniques are not given. This is one of the primary stages which differentiates competing photogrammetry software, and is clearly treated as a valuable and confidential part of these algorithms.

A PRACTICAL ALGORITHM FOR SMALL-AREA SOFTCOPY PHOTOGRAMMETRY

Motivation

For the calculation of absolute height values from large parallax maps, ground control points with precise and accurate *x,y,z* co-ordinates are required, in order to determine the tilt, swing, and azimuth angles which form part of the exterior orientation system of a photograph. Accurate lens calibration models are also required, for calculation of the interior orientation system of a photograph. However, due to the fact that the photographs used by archaeologists are often borrowed rather than personally commissioned, this information is typically not available. There is clearly a useful role for a DEM

generation system that does not require precise ground control or technical information about the camera.

Simple equations that disregard interior and exterior orientations and other sources of geometric image distortion, can be used to compute heights from parallax maps. However, as a general rule, if two points are located more than one inch apart on a photograph, errors in calculated elevation differences, due primarily to camera tilt and paper shrinkage, render the calculations unreliable (Wong 1980). It is clear therefore that, for a system that does not incorporate complex image model information, DEMs which cover only small areas of a photograph (less than 1 square inch) can be accurately created. This would effectively allow DEMs of individual monuments and their immediate vicinities to be created.

The Algorithm

Figure 5 presents the steps involved in the algorithm under discussion. Each of the steps is further discussed is the following sections.

---

1. Definition of search region
  1.1 User identification of 3 or more corresponding tie points in the two images
  1.2 Calculation of relative search region from tie points
2. Performance of first cross-correlation pass
  2.1 Preparation of photographs using edge emphasis filter or histogram equalisation
  2.2 Calculation of parallax values at each pixel in left image
    2.2.1 Further dramatic reduction of search region to centre on the pixel in the right image suggested by the modal parallax of the nearest 12 correlated neighbours, if the frequency of this value is at least 4
    2.2.2 Calculation of cross-correlation coefficient for each position in the search region and storage of the x and y parallax values that produced this lowest coefficient
3. Performance of second cross-correlation pass. The same as step 2, except certain parameters may be different
4. Detection and removal of obstacles, errors and noise
  4.1 Performance of 'flood-rejection' operation
  4.2 Estimation of rejected pixels from non-rejected neighbours
5. Calculation of x/y and z scales of DEM, using knowledge of photograph scales, scan resolutions, and flying height of the aeroplane

---

Figure 5: The DEM generation algorithm proposed for small-area topographic modelling. The information required for step 5 of the process is supplied by the user (flying height; scan resolution) and calculated from user-supplied approximate ground control points (principal point locations; photograph scales and orientations).

Definition of the Search Region

This is a crucial step in digital photogrammetry systems for two reasons: in order to speed up the cross correlation process, and in order to reduce the number of incorrect matches made. The technique used by this algorithm makes use of 'tie-point' information from the user. This involves the user identifying a small number of points in the first image, and the corresponding points in the second. From this, the range of x and y direction parallax values that need to be tested, for each pixel, are calculated. In practice, it is necessary to enlarge this range, to take into account the fact that the user may not have identified the highest and lowest points in the landscape. Reducing the search region in this way clearly improves the speed and accuracy of cross-correlation (figure 6).

Figure 6   Cross correlation with tie-point derived search region reduction. (Top) Stereo images of an archaeological monument (Rathcroghan mound, Co. Roscommon), scanned at 600 d.p.i., consisting of approx. 90,000 pixels. The results of applying 7x7 pixel correlation matching: (bottom left) undirected - calculation time 94 minutes on a Pentium 166. Apart from pixels around the base of the monument, where there is good evidence for correlation matching, most of this parallax map is simply noise; (bottom right) tie point directed - calculation time 4 minutes. The majority of pixels have now been correctly matched, though considerable amounts of noise are still evident.

A further, more dramatic region narrowing approach is applied during the cross-correlation stages (described below): this allows smaller regions to be defined from local pixel information, thereby providing greater improvements.

Two Cross-Correlation Stages

The bulk of the computational work is carried out during the cross-correlation stages, where each of the pixels in the first image is matched to a corresponding pixel in the second image. Further reduction of search regions is effected on a pixel-by-pixel basis, using parallax information from the neighbourhood of the pixel of current interest. This 'neighbour-assisted' search narrowing is proposed as an alternative to the more typical coarse-to-fine approach. The principal is, if several neighbours of a pixel have had identical parallax values already calculated, then this provides evidence for the dramatic reduction of the search region, to a small window around the pixel at the position suggested by these neighbours.

Two entire passes are made through the images during cross correlation: the output from the first stage is used to narrow search regions in the second stage. The second pass consists of the same operations as the first, except that some or all of the following parameters will be different:

- The frequency of neighbours' modal parallax required for dramatic search region reduction. On the first pass, at least 4 neighbours of the closest 12 must have identical parallax values for search region reduction to proceed, while on the second pass, at least 8 of the closest 24 must agree;
- The size of cross-correlation window used, in order to take advantage of different window sizes. The size in pixels of a side of the window on the first pass is suggested as $2*R/200+5$, where R is the scan resolution of the photos, in dots per inch. The default size of a side of the window on the second pass is $2*R/200+1$, with a minimum value of 5.

When performing cross-correlation, the size of the window used is of vital importance. Large windows are relatively insensitive to noise, though they are expensive in terms of computation time, and they destroy small topographic features. Small windows, though sensitive to noise, require less computation and allow accurate matching co-ordinates to be calculated, thereby preserving small features. Figure 7 illustrates the effect of large and small correlation windows.

This neighbour-assisted approach allows information from the first pass to assist correlation in the second pass, by cutting down the search region to small areas centring on hypothesised matching pixels. The size of the correlation window in the second pass may be as small as 5x5 pixels, allowing the accurate extraction of small topographic features without excessive amounts of erroneous matches. In practice, the factor which limits the size of the window in the first pass is computation time.

Figure 7   The effect of different correlation window sizes on DEM production using stereo images of the ringfort of figure 3. (Left) The result when using a 23x23 correlation window; (middle) using a 5x5 correlation window; (right) using two correlation passes, the first with window size 23x23, and the second with size 5x5.

The algorithm makes use of an approach to improving image texture which is demonstrably more effective than histogram equalisation in many cases. This is to apply an edge enhancing filter, such as the Laplacian (Marr & Hildreth 1980) - which is essentially a measure of the dissimilarity between a pixel and its neighbours. This greatly reduces the number of cross correlation errors, particularly in areas of poor texture, by emphasising high frequency features, such as small clumps of grass, which are visible in both images of a stereo pair (figure 8).


Figure 8   Texture improvement using edge detection. (Left) Edge enhanced 600dpi ringfort image; (right) DEM produced from a single correlation pass with a 9x9 window. This compares favourably to the result shown in figure 4.

Detection and Removal of Obstacles, Blunders, and Noise

Given a dense parallax map, such as those produced by the current approach, it is proposed that regions of noise due to obstacles or blunders may be identified because of their relatively small areas. A parallax map can therefore by partitioned into connected regions of homogenous parallax values, and any regions with areas below a given threshold can be rejected and replaced with interpolated values. In order to calculate the area of a

connected region of identical parallax values, an operation similar to the 'flood fill' of graphics software is performed (see pseudocode in figure 9). Figure 10 illustrates the effectiveness of this technique: even relatively large patches of noise are successfully detected and removed, without blurring or affecting in any way the shape of acceptable contours.

```
Given: 2-Dimensional parallax map P, minimum acceptable patch
       area S

Procedure Reject          /* controlling procedure */
       FOR each pixel (x,y) in P where P(x,y) > 0:
         Call FloodFix(x,y,0)
          IF a value of less than S was returned, THEN reject all
           pixels stored during the recursive FloodFix calls
       FOR each non-rejected pixel (x,y) in P:
         Set P(x,y) to abs(P(x,y))
End Procedure

                 /* recursive region size calculation function */
Function FloodReject(x,y,counter)
       LET c = P(x,y)
       Set P(x,y) to -P(x,y), so that it is not visited again
       FOR each (x₁,y₁) where x₁={x-1,x,x+1}, y₁={y-1,y,y+1}:
           IF P(x₁,y₁)=c THEN recursively call
       FloodFix(x₁,y₁,counter)
       Increment counter    /* size of current region so far */
       IF counter < S THEN store (x,y)
Return counter
```

Figure 9   Pseudo code implementation of an efficient recursive "flood-reject" operation, which identifies and rejects small regions of area less than S pixels in a parallax map. All parallax values are assumed to be positive at the start. By setting them negative, a situation whereby neighbouring pixels repeatedly flood into each other is avoided: each pixel is visited exactly once.

Calculation of x/y and z Scales

The distance on the ground represented by a distance of 1 pixel in a scanned photograph is estimated through the use of user-supplied control points, extracted for example from Ordnance Survey maps. Given a number of control points, the meaning of 1 pixel implied by each possible pair of control points is determined, and the average is calculated. This provides results of acceptable accuracy.

Figure 10   The application of the 'flood-rejection' operation to a larger section of a photograph. (Top) Original image from a stereo pair; (middle) 'raw' DEM with errors due to obstacles and blunders apparent; (bottom) final DEM, following error removal.

The height (*z*) scale of a DEM, i.e. the meaning in metres of elevation represented by 1 measured parallax unit, is also required. Due primarily to low-frequency geometric distortions in aerial photographs, an empirical approach similar to that used to estimate the distance scale was found to be wildly unreliable. The current algorithm therefore calculates the height scale theoretically from an idealised geometric object and image model (figure 11). With all units of measurement expressed in terms of metres, the height in metres that is represented by 1 unit of parallax can be calculated from the airbase (distance travelled between exposures) and average flying height above the landscape. The airbase is calculated as the distance in pixels between conjugate principal points[1], multiplied by the distance on the ground represented by each pixel. Flying height information, which is automatically recorded on aerial survey photographs, is supplied by the user.


Figure 11   Object and image model, allowing the calculation of units of parallax.

Referring to the object and image model of figure 11, by similar triangles:

$$B/p = (H-h)/h$$
$$\Rightarrow \quad hB = pH - ph$$
$$\Rightarrow \quad h[B+p] = pH$$
$$\Rightarrow \quad h = (pH)/(B+p)$$

Therefore, if   **∆p = 1 pixel**

---

[1] Given a stereo pair of photographs with good overlap, the principal (centre) point of each photograph - which is approximately the ground location of the aeroplane at the time of exposure - will also be visible on the other photograph. The distance between these points provides the airbase.

Then $$Dh = (pixdist*H)/(B+pixdist) \text{ metres}$$

where *pixdist* is the ground distance represented by one pixel in the images[2].

ASSESSMENT OF RESULTS

Overall Accuracy

For an assessment of the accuracy of the algorithm presented in this paper, see figure 12, which provides a visual comparison between (a) an EDM-derived DEM of Rathcroghan mound, and (b) a DEM of the same monument generated using this algorithm. For statistical testing, Pearson's correlation coefficient was calculated on a line-by-line (horizontal transept) basis, between the two DEMs. The average coefficient for these lines was 0.96.

In figure 13, the heights in the two DEMs of evenly sampled pixels are plotted against each other: the match between heights is very close, though not identical. The error bias of the photogrammetry model is +8 cm. per pixel, and the standard deviation of errors is 34 cm.: therefore the overall accuracy is to about 42 cm., which is approximately 1/2,500 of flying height above the landscape[3]. This is better than that reported by some commercial photogrammetry software (Stallmann 1995). However, the current approach can only obtain this level of accuracy for small areas of the full overlap between photographs, in which low frequency geometric distortion is not an issue. It is presumed that the reported accuracy of commercial techniques is for full stereo overlaps.

Figure 12   DEMs of Rathcroghan *mound, (left) generated by several weeks of ground survey, and (right) through use of the current algorithm. The* x/y *resolution of the second DEM was reduced, and the* z *resolution of the first DEM was also reduced, to make these images visually compatible.*

---

[2] Similar derivations can be found in Wong (1980) and Ritchie *et al.* (1988).

[3] The flying height above the landscape was approximately 1 km.

Figure 13   Scatter plot of heights in the two DEMs of figure 12.

EDM-derived data for another monument from the Rathcroghan complex, the enclosure called *Rath na Darve*, was also obtained. Figure 14 illustrates this DEM alongside one created using the algorithms described in this paper. In this test, the error bias of the photogrammetry model was +11 cm. per pixel, and the standard deviation of errors was 58 cm.: therefore the overall accuracy is to about 69 cm., which is approximately 1/1,500 of flying height above the landscape. As before, Pearson's correlation coefficient was calculated on a line-by-line basis, between the two DEMs. The average coefficient for these lines was 0.94.


Figure 14   DEMs of Rath na Darve*. (Left) The monument as it appears in an aerial photograph; (middle) DEM generated using ground survey, and (right) using the current algorithm.*

Neighbour Assisted Multiple-Pass Cross-Correlation

The neighbour-assisted technique for multiple pass cross-correlation is an alternative to the multiple resolution coarse-to-fine approach, which is essentially also a neighbourhood-assisted approach, since a single pixel at a coarse resolution represents a square region of pixels at a finer resolution. The primary improvement represented by the neighbour-assisted technique developed in the current research is that several, rather than one, cross-correlation operation of a pass are used to direct an operation on the next

pass. This allows a degree of error tolerance, or noise insensitivity, to be introduced. The neighbour assisted approach directs (i.e. region narrows) cross-correlation using the modal value of a pixel's neighbours, and furthermore, it only performs this region narrowing if this modal value has a frequency of greater than a given threshold. By comparison, the coarse-to-fine approach is equivalent to applying region narrowing based on neighbourhood means, and with no facility for determining whether there is enough evidence to perform this narrowing for a particular pixel.

In order to quantitatively compare neighbour assisted multiple pass cross-correlation with coarse-to-fine cross-correlation, a number of test DEMs were created, and the 'flood rejection' operation used in order to determine approximately the number of erroneous pixels in each. In addition, the computation times were noted. Histogram equalisation and control-point based search region narrowing were applied in all cases. The same correlation window sizes, and the same size narrowed search region were also used throughout.

Figure 15 presents one of the comparisons between DEMs generated by the standard and new techniques. It is evident that the new technique provides superior noise tolerance when dealing with poor texture images and relatively flat ground. Another test involved relatively steep ground, and in this case the traditional technique proved to be more computationally efficient, though not more accurate.

Figure 15   Comparison of DEMs made from 600dpi scans of Rath na Darve. *Rejected pixels have been coloured black rather than interpolated. (Left) Coarse-to-fine DEM, which took 6 mins 5 secs, and yielded 4998 error pixels (7.9% of the DEM). (Right) Neighbour-assisted DEM, which took 6 mins 20 secs, and yielded 4666 error pixels (6.0% of the DEM). 13996 pixels were lost on the right and top of the coarse-to-fine model, and additional pixels on left and bottom edges still need to be removed.*

Texture Improvement

The edge enhancement of images can be used as a technique to improve texture, and therefore assist cross-correlation accuracy and speed. Though a

simple operation, this has not been previously published as an alternative to histogram equalisation. In order to quantitatively assess the suitability of edge enhancement as a preparation technique for stereo images before cross-correlation, test DEMs were created using as input raw stereo images, histogram equalised images, and edge enhanced images. As before, performances were rated in terms of the amount of pixels rejected by a 'flood reject' operation, and the overall computation time. It is clear that raw landscape images are not generally suitable for cross-correlation, since in the images with poor texture and reasonable texture, disastrously large rates of error occurred. Figure 16 presents one of the test DEMs.


Figure 16   Comparison of image preparation techniques using images of a ringfort with low texture. (Left) DEM from raw images, which took 1 min 55 secs, and yielded 12245 rejected pixels (58.4%) ; (middle) DEM from histogram equalised images, which took 1 min 30secs, and yielded 2705 (12.9%) rejected pixels; (right) DEM from edge enhanced images, which took 1 min 35 secs, and yielded 1295 (6.2%) rejected pixels.


Figure 17   Comparison of the rates of cross-correlation error in DEMs generated from differently prepared images.

In the tests using images with good texture and reasonable texture, there was little significant difference between the results of the histogram equalised and edge enhanced approaches, though the latter performed better in both cases. The superiority of the edge enhanced approach is most evident in the test which used low texture images (figure 16). The traditional image preparation approach, which essentially re-maps brightness values, does not fully solve the problem of poor texture. Edge enhancement, however, fundamentally changes images, into first or second spatial-derivative maps of brightness change. This has the effect of emphasising high frequency tonal changes, while maintaining the morphology of an image. Figure 17 presents the results of these error tests in graphical form. The texture of a landscape image has little effect on edge enhanced preparation.

CONCLUSIONS

The algorithm presented in this paper meets the specific topographic modelling requirements of archaeologists, which include good accuracy, simplicity, labour efficiency and minimal requirements for additional information about the photographs or camera. The small-area DEMs that are produced by the technique are suited to modelling individual monuments, which can be used to automatically derive other archaeologically relevant information, such as slope and aspect (see Redfern 1997; Redfern *et al.* 1998).

Although softcopy topographic photogrammetry can recently be termed a 'solved problem', there are aspects of the algorithm described in this paper that represent improvements on the existing published approaches. The primary innovative aspects of the technique are:

- The neighbour-assisted multiple pass cross-correlation technique, as an alternative to the standard coarse-to-fine image pyramid approach. This new approach provides superior error tolerance, though is marginally less computationally efficient than the traditional one;
- The application of edge enhancement to images in order to improve texture and therefore cross-correlation speed and accuracy, as an alternative to the standard histogram equalisation. Edge enhancement provides significant improvements in images or regions of an image where texture is poor, while performing similarly to histogram equalisation in images where texture is good;
- The identification of obstacles, cross-correlation blunders and other DEM noise, through the application of a recursive 'flood-reject' function, which partitions a DEM into connected regions of homogenous height value while calculating the areas of these regions. Noise, identified as small regions, is removed without affecting in any way the shape or size of

acceptable areas. It has been noted by several researchers that robust DEM noise correction is the primary aspect of softcopy photogrammetry still requiring further research.

ACKNOWLEDGEMENT

REFERENCES

Ackermann, F. and P. Krzystek. 1995. New Investigations into the Technical Performance of Automatic DEM Generation. *Proceedings ASPRS/ACSM Annual Convention,* **2**. 488-500.

Ackermann, F. 1996. Some Considerations About Feature Matching for the Automatic Generation of Digital Elevation Models. *OEEPE Workshop on the Application of Digital Photogrammetric Workstations.*

Al-Rajeh, Khalid M. 1994. *Computer-Aided Matching of a Sequence of Images*. Ph.D. Thesis, University of Bradford, UK.

Dupéret, Alain. 1996. Automatic Derivation of a DTM to Produce Contour Lines. *OEEPE Workshop on the Application of Digital Photogrammetric Workstations.*

Goshtasby, A., S.H. Gage, and J.F. Bartholic. 1984. A Two-Stage Cross Correlation Approach to Template Matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **6**. 374-8.

Haigh, JGB. 1993. Practical Experience in Creating Digital Terrain Models. *In* Andresen, Madsen, and Scollar (ed). *CAA92 - Computing the Past: Computer Applications and Quantitative Methods in Archaeology.* Aarhus University Press: Aarhus. 67-74.

Heipke, C. 1996. Overview of Image Matching Techniques. *OEEPE Workshop on the Application of Digital Photogrammetric Workstations.*

Kirsch, RA. 1997. Photogrammetric Reconstruction of Petroglyphs. *American Indian Rock Art*, **23**. 177-82.

Krzystek, Peter. 1995. Generation of Digital Elevation Models. *Second Course in Digital Photogrammetry.* Institut für Photogrammetrie, Universität Bonn.

Kvamme, KL. 1990. GIS Algorithms and Their Effects on Regional Archaeological Analysis. *In* Allen, Kathleen M., Green, S., Zubrow, E. (ed) *Interpreting Space: GIS and Archaeology.* Taylor & Francis: London. 112-25.

Marr, D. and E. Hildreth. 1980. Theory of Edge Detection. *Proceedings of the Royal Society London*, **207**. 187-217.

Ritchie, William, Michael Wood, Robert Wright and David Tait. 1988. *Surveying and Mapping for Field Scientists*. Longman Scientific & Technical: Harlow, Essex, UK.

Redfern, Sam. 1997. Computer Assisted Classification from Aerial Photographs. *AARGnews* **14**. 33-8.

Redfern, Sam, Gerard Lyons and R. Michael Redfern. 1998 (forthcoming). The Automatic Morphological Description and Classification of Archaeological Monuments from Vertical Aerial Photographs. *Proceedings of OEMI/IMVIP Joint Conference*, Maynooth, Ireland, 1998.

Slama, Chester, Charles Theurer and Soren Henriksen (ed.) 1980. *Manual of Photogrammetry*, 4th edition. American Society of Photogrammetry: Va.

Stallmann, Dirk. 1995. *Automatic Digital Terrain Model Generation*. [Online] available http://www.p.igp.ethz.ch/p02/projects/autodtm/p1.html, March 20th 1998.

Wong, KW. 1980. Basic Mathematics of Photogrammetry. *In* Chester Slama, Charles Theurer and Soren Henriksen (ed.) *Manual of Photogrammetry*, 4th edition. American Society of Photogrammetry: Va. 37-101.