Jon Bradley
Mike Fletcher

Extraction and visualisation of information from ground penetrating radar surveys

1 Introduction

The analysis of ground penetrating radar (GPR) data is, for a variety of reasons, an imprecise and time consuming activity. Manual evaluation of the data involves a high level of experience and takes an inordinate amount of time. It often requires that subjective decisions be made by the analyst, and this leads to the introduction of bias (both systematic and non-systematic) into the final results.

Computer analysis of GPR data has, until now, tended to focus either on the use of complex signal processing algorithms in an attempt to remove all sources of distortion from the image (Daniels *et al*. 1988: 298-303), or on simple image enhancement techniques to make the image easier for the analyst to understand (Blake 1995; Daniels *et al*. 1988: 297-298). The former approach, whilst effective in simple scenarios, tends to fall down on more complex sites, whilst the latter fails to remove the subjectivity from the process.

The objective of this paper is to outline and demonstrate an alternative, statistically based, approach to the extraction of information from GPR data. Such an approach requires some measure of the level of *radar activity* at each point of a radar image. Two possible solutions are discussed in this paper — *the sum of squared errors (SSE)* and the *k* measures.

The basic implementation of the technique yields a plan view of the site, which is useful as a first stage of interpretation. The drawback is that it contains no depth information, but a simple elaboration of the analysis allows rudimentary depth information to be extracted from the data in the form of horizontal 'slices' through the site.

In order to illustrate the practical application of the techniques outlined in this paper, two case studies are examined. They are based on survey work done at Worcester Cathedral and at a Bronze Age ring cairn at Stapeley Hill near Welshpool (Stratascan 1994).

2 Activity measure analysis

An analyst might manually examine radar data with the intention of picking out features of interest. The criteria which decide whether a feature is of interest may be very varied. Factors influencing the choice include contrast, coherency, size and shape of features. The interest/no interest decision is obviously a subjective one, hence the need for a high level of expertise, and any way of introducing objectivity into the process would therefore be of benefit in terms of reliability of results. A simple way of addressing this problem is the use of activity measures. Their function is to provide an objective, quantitive assessment of the degree of interest in a chosen region of the radar image. The criteria on which this value is based varies with the particular measure chosen. This paper discusses two of the simpler possibilities.

2.1 Activity measures

Two statistically based activity measures were developed for use with the technique.

The *sum of squared errors (SSE)* measure quantifies the degree of deviation away from the mean intensity of the image. It is given by

$$SSE = \sum_{i=1}^{n} (x_i - \bar{x})^2 \qquad (1)$$

where x_i is the intensity of the ith pixel, $\bar{x}$ is the mean pixel intensity and n is the number of pixels under consideration.

The k measure combines the SSE measure discussed above with second order information on the abruptness of change in the image. It is normalised between 0 and 1 and is given by

$$k = 4SSE/naC \qquad (2)$$

where n is the number of pixels under consideration, a is the sum of the pixel intensities and C (total change) is given by

$$C = \sum_{i=1}^{n-1} |x_i - x_{i+1}| \qquad (3)$$

2.2 Image preprocessing

Before analysis of the radar survey can take place certain preprocessing steps have to be taken. Firstly the images must be cropped to remove any scrap data caused by the antenna unit standing still at either end of the transect (fig. 1). Then the image is scanned to establish the positions of the

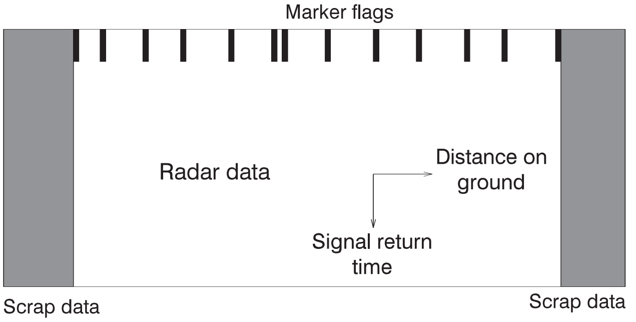

Figure 1. Features of a typical radar image.

ground distance marker flags. These correspond to intervals of one metre on the ground and are inserted manually by the antenna operator. Because the antenna unit is dragged at a non-constant rate (particularly on rough terrain), the markers in the image are generally unevenly spaced.
In order to allow an accurate plan of the survey to be constructed, these irregularities must be removed prior to processing. This is done using simple linear interpolation techniques to rescale all intervals to a standard, pre-determined width. Once the horizontal rectification process has been performed on all of the images, they will be of a standard horizontal scale, and the analysis can continue.

2.3 Activity profiling, mapping and visualisation

The objective of the analysis is to reduce each survey transect (a two-dimensional image) to a single string of values — *an activity profile*. Each value in what is now a one-dimensional string of data represents the level of activity below that point of the transect *for the full depth of the survey*. Activity values are calculated at predetermined intervals along the transects and once profiles have been generated for each transect, the data can be used to build a plan of the activity over the survey area.

The surveying in both of the case studies was carried out along orthogonal transects to produce grids of squares. Three approaches to visualising the activity of the survey area were tested. A problem arises here in deciding which of the two possible activity values at transect intersections to use. Differences in the measured activity values for differently oriented transects could arise for various reasons including calibration drift, asymmetries of the radar beam geometry and random effects caused by terrain and operator error. So far no work has been done to investigate these problems quantitatively. For the sake of simplicity the approach adopted in this paper was to use the higher of the two activity values.

The first method, *the activity grid*, represents each transects activity profile as a grey level strip arranged in a grid on a black background. An example of this technique is shown in figure 2a. Although useful as a preliminary visualisation tool, this approach proves unsatisfactory in one important respect; the use of linear grey level strips on a monotone background can sometimes deceive the eye into seeing spurious linear features in the data.

This potentially serious drawback is overcome by the use of a *simple activity map*, an example of which is given in figure 2b. Here the data is presented as a colour coded map (with the option of contours or pseudo-surface representation). To achieve this a second order surface is fitted to the square grid of points formed by the transect intersection points. The drawback of this approach is that it makes minimal use of the available information. The sampling rate along survey transects is extremely high, but this approach uses only a tiny fraction of the sampled points.

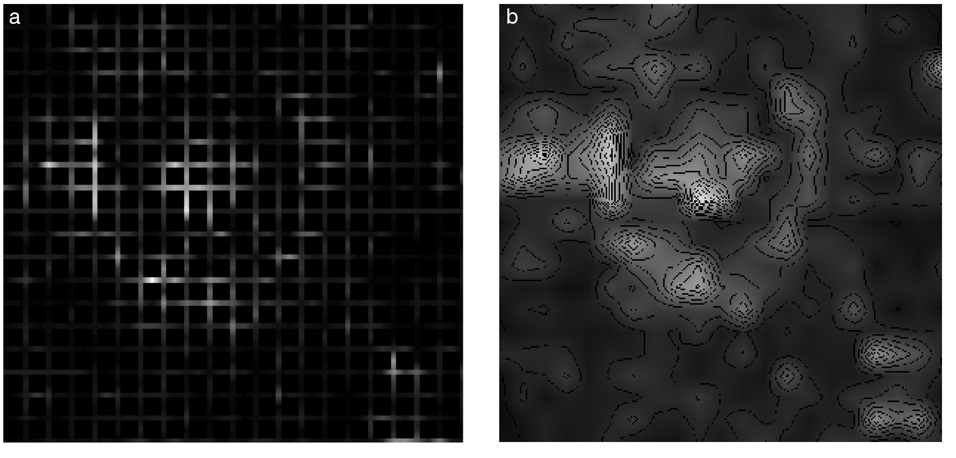

Figure 2. An activity grid and the corresponding simple activity map.

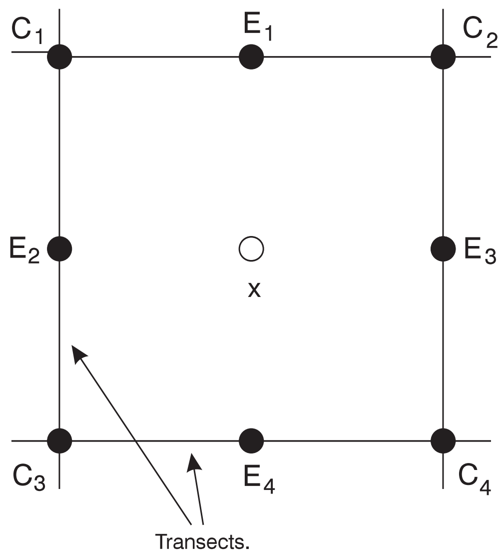

Figure 3. The scale-doubling interpolation algorithm.

The third and more sophisticated approach, *the interpolated activity map*, goes some way to addressing the problem of utilization of data. The interpolation is a two stage process. The basis for the first stage (which effectively doubles the scale) is shown in figure 3 and equation (4)

$$x = 0.5 \sum_{i=1}^{4} E_i - 0.25 \sum_{i=1}^{4} C_i \quad (4)$$

where x is the unknown value, C is the activity value at each transect intersection and E is the activity value midway between transect intersections. This interpolation step produces a square grid at twice the resolution to that used by the simple activity map. The second interpolation step is then identical to that in the previous technique, but produces a map on double the scale.

2.4 Case studies

The following two case studies are used to illustrate the potential of the new technique. For the sake of brevity only the *interpolated* results of the *SSE* analysis are presented. The *k* analysis appears, on the basis of these two studies, to produce similar results with only subtle differences in contrast between the two.

2.4.1 Stapeley Hill ring cairn

Stapeley Hill (OS Ref. SO 313 991) is a small ring cairn of supposed Bronze Age origin. Topographical surveying carried out on the site (Fletcher/Spicer 1990; Spicer 1991), has confirmed the presence of ridge and furrow and field walls. A lit three-dimensional surface representation of the topography of the site from a more recent survey (Stratascan 1994) is shown in figure 4. The vertical heights are exaggerated by a factor of three to emphasise features. Note the edge of a wide ridge in the foreground of the image caused by the remains of an ancient field wall, and behind it the obvious ring structure with a small central mound. Ridge and furrow are also in evidence to the left of the ring cairn (although not so obviously). The scale of the site is 20 m square, north is to the top left.

Radar surveying was carried out at the site in May 1994 using a 500 MHz antenna on a two directional 1 m grid. An

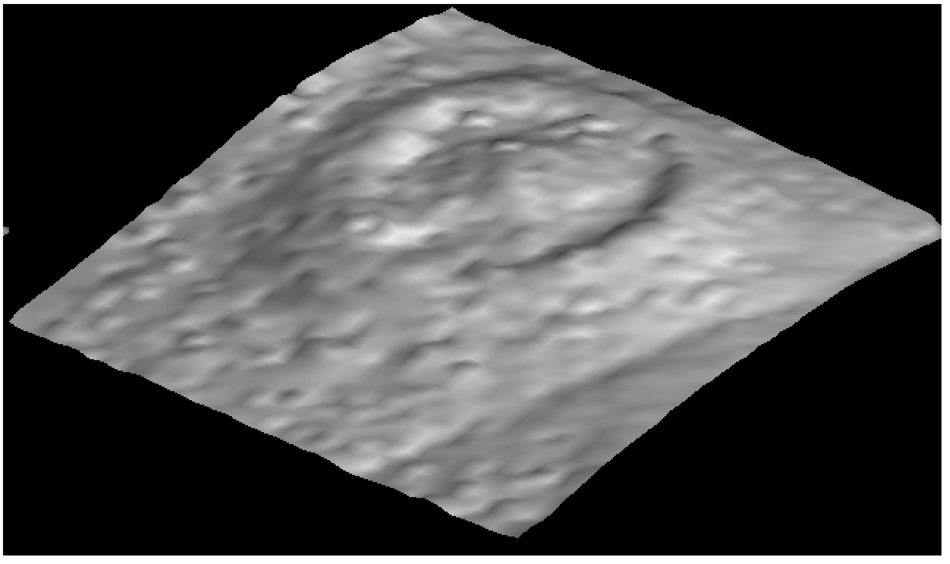

Figure 4. Stapeley Hill - topography.

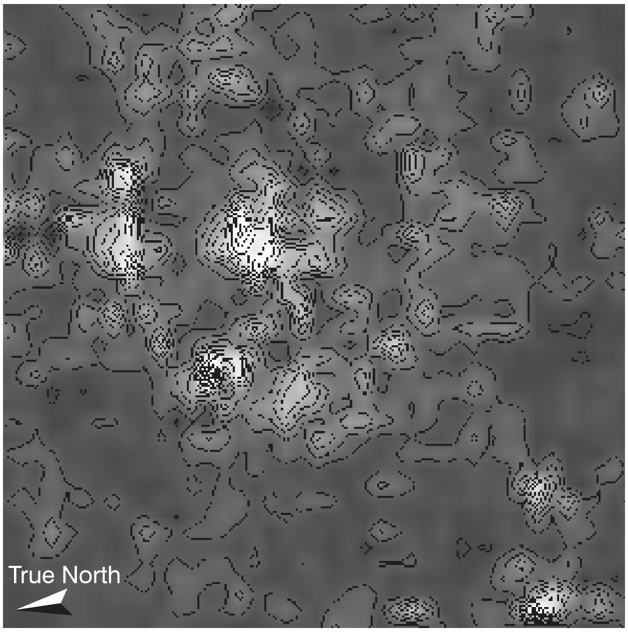

Figure 5. Stapeley Hill - radar activity (SSE).

interpolated activity map of the SSE values for the site is shown in figure 5.

The activity map reveals a rough ring shaped area of enhanced radar activity (top centre of image) with an irregular central feature. A further area of high activity lies at the bottom right of the image and is probably associated with field wall remains. Superimposition of the radar activity map on the topography (fig. 6) reveals good correspondences between radar and topographical evidence for field wall remains, between the positions of the circular bank of the ring work and the circular feature on the radar map, and between the central mound and the central radar feature.

This evidence probably points to the presence of a large proportion of stony material in the circular bank. The activity strength of the bank is weaker on the SE side of the ring. This may be due to the topography of the site (fig. 7) or to possible removal of stone associated with the ridge and furrow or with the field wall. The greater strength and degree of coherency of the central feature may point to the presence of a large, coherent mass of stone — possibly a cist.

2.4.2 Worcester Cathedral

The second case study is based on a survey conducted at Worcester Cathedral which was expected to reveal traces of foundations from a previous structure. The survey covered a rectangular area of approximately 280 square metres and was done on a regular 1 m grid in two directions, again at 500 MHz. The SSE analysis results are shown in figure 8.

The results are not as revealing as those for Stapeley Hill. The most obvious feature is the area of high activity in the top left of the survey. This feature is possibly the high activity end of a linear region of enhanced activity extending from top left to bottom right of the map, though this is by no means certain. This may correspond to the path of an old drainage channel which still retains some moisture. Some regular features can also be discerned at bottom centre and bottom left, and it is tempting to associate these with building foundations.

3 Retrieval of depth information

3.1 Simple stratification

The approach as it stands at this point yields only two-dimensional information in the form of a plan view of the radar activity levels over the survey site. This should, in many cases, provide a coarse indicator of the *presence* of features of interest, but it may be that a more comprehensive, three-dimensional picture of the site is required.

Accurate extraction of three-dimensional information from radar data is a problematic process. Various factors including surface topography and lateral inhomogeneities in the electro-magnetic properties of the soil matrix cause vertical distortions in the data which cannot be easily rectified. A first step towards the extraction of depth information from a survey is to simply ‘stratify’ the radar images and to analyse each layer separately. It should be emphasised that vertical distance on the radar images does not correspond directly to vertical distance below the ground surface but to the signal return time of the radar emissions. This means that what is produced by this process is a series of activity maps which, *approximately* speaking, depict the activity at different distances below the ground surface. Theoretically, the simpler the site, the better this approximation will be.

This approach is similar in nature to that taken by Milligan and Aitkin (Milligan/Aitken 1993: 26-39), but with several distinct advantages. The approach adopted by Milligan and Aitken is limited in scope in that the horizontal (signal return time) slices are produced by direct extraction of pixel values from a *uni-directional* set of parallel radar images. No ‘interest’ values such as the activity values developed herein were employed and no interpolation was attempted since interpolation from a uni-directional survey tended to result in spurious linear features at right angles to the transect direction.

3.2 Case studies

3.2.1 Stapeley Hill

The results of applying the stratification process to Stapeley Hill are interesting. Examination of the images in figure 9 shows that the NE side of the ring shows up clearly only on

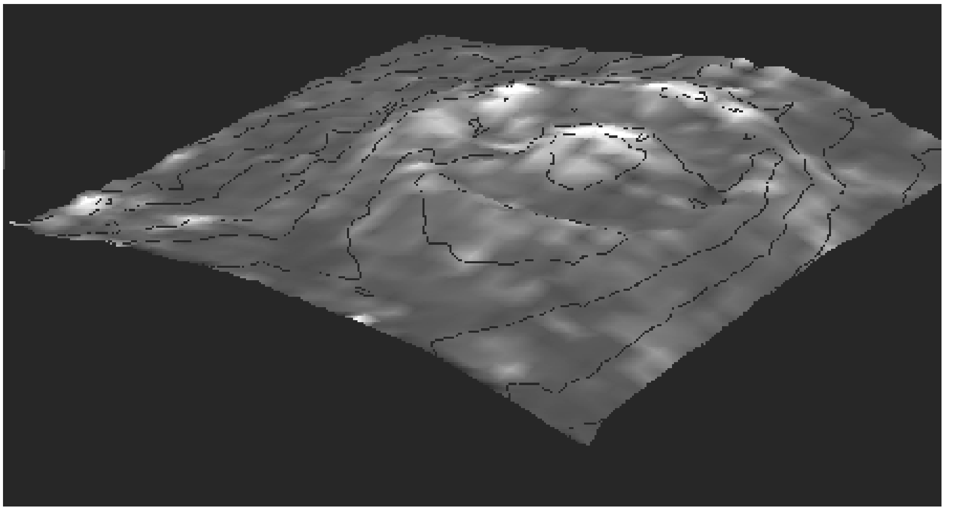

Figure 6. Stapeley Hill - radar activity superimposed on topography.

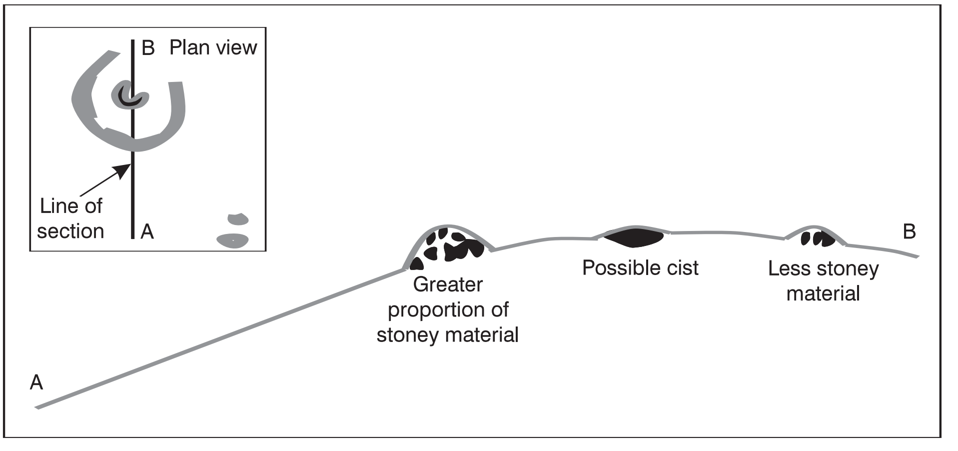

Figure 7. One possible explanation for the discontinuous nature of the radar activity in the bank of the ring cairn.

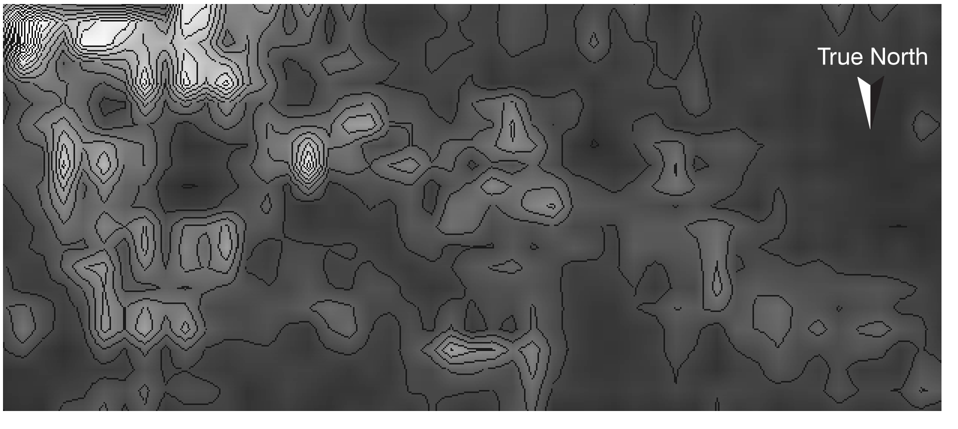

Figure 8. Worcester Cathedral - SSE results.

the first (shallowest) map, whilst the rest of the structure shows up on all of the maps. This lends support to the conclusions that the ring is more substantially constructed on the down slope (NW and W) sides. The central structure does not show up at all in (a), appears very strongly in (b), fades away again in (c), then makes a reappearance in (d). This lends support to the hypothesised presence of a structured central feature i.e. a cist. A further feature of

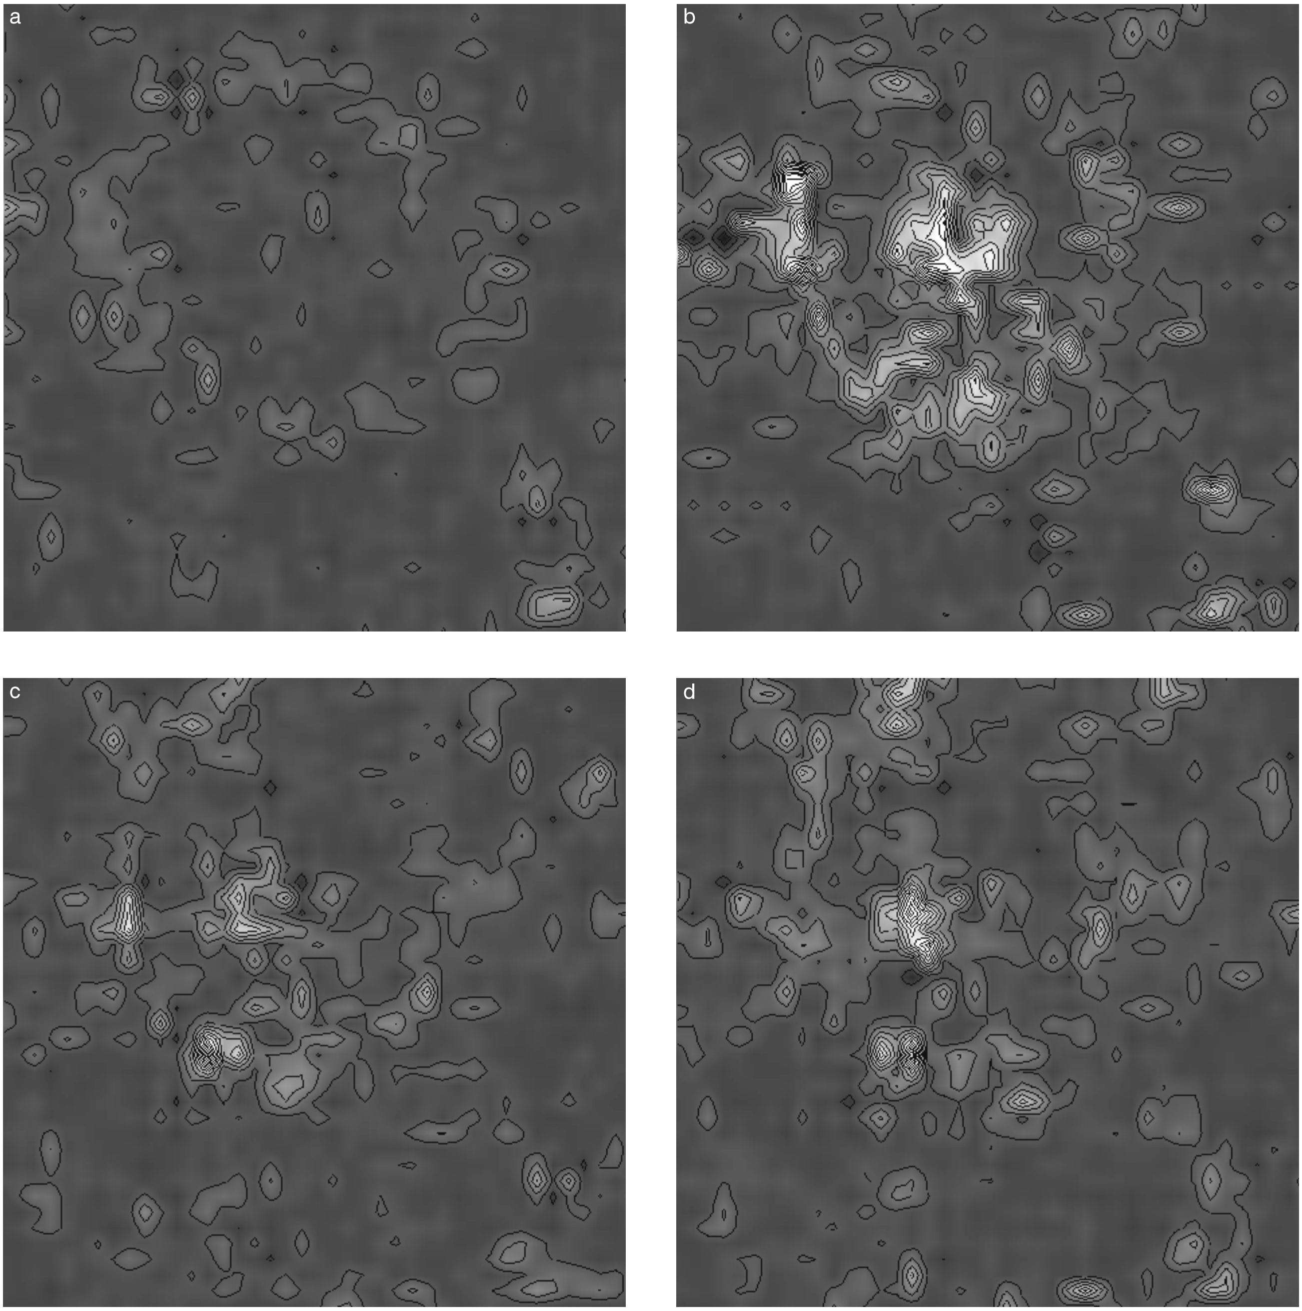

Figure 9. Four consecutive SSE slices through Stapeley Hill.

possible interest is a large anomaly which shows up very strongly in the NW side of the ring at greater depths (in maps (c) and (d)). The radar signatures associated with the field wall appear concentrated mainly at intermediate to shallow depths (map (b)). This would appear consistent with the supposed later origins of the wall.

3.2.2 Worcester Cathedral

The use of the simple stratification method on the Worcester Cathedral data is very revealing. The sequence of maps (fig. 10) going from (a) the shallowest, to (e) the deepest, clearly confirm the presence of the suspected linear feature. This provides a good example of how problems caused by overlapping features in successive strata can be easily and effectively overcome.

4 Further work

4.1 Image filtering and depth information

In order to progress further certain issues need to be addressed. The results obtained thus far, although interesting, are flawed in that the presence of detectable objects in the upper layers of the site tends to affect the activity levels calculated for areas of the radar image which lie below. This is due to the fact that the radar pulse produced by the apparatus is not instantaneous in nature, but consists of a decaying wave train. This leads to a periodic *echoing* effect where 'ghosts' of objects appear in the radar image below the main object signature.

Because these effects are periodic, they should prove amenable to manipulation by frequency domain techniques such as Fourier analysis. The main problem here is that the removal of the unwanted 'ghosting' is not simply a matter of detecting the echo frequency and filtering it from the image. Although this would have the desired effect, it would also mean that the primary signal is filtered out. Further investigation of the problem is obviously necessary.

4.2 Improving data utilization

Data points are available at very small intervals along the transects (so much so in fact, that the data can almost be considered to be continuous). Some attempt has been made to use interpolation techniques to make better use of the available data (section 2.3). In spite of this there is still a great deal of room for improvement, perhaps by treating the data as irregular xyz values and using routines which can deal with this type of data.

5 Conclusions

This paper shows that it is possible to produce simple, quick and, most importantly, objective routines for analysing data from ground penetrating radar surveys. Two statistical measures of radar activity — the SSE and k values — have been developed. Other measures, possibly relying on frequency domain information, have yet to be examined.

A simple extension of the basic technique allows some appreciation of the three-dimensional structure of the survey site to be obtained. As has been demonstrated, this often allows extra inferences and conclusions about the site to be drawn. An initial evaluation detailed in the case studies examined above, indicates that this technique may prove to be an extremely useful tool, but much work still needs to be done in this area in order to iron out difficulties caused by the radar ghosting effect detailed in section 4.1.

It is not possible to say at this stage which measure of activity produces the better results, and work using a known 'model' site is necessary in order to test them.

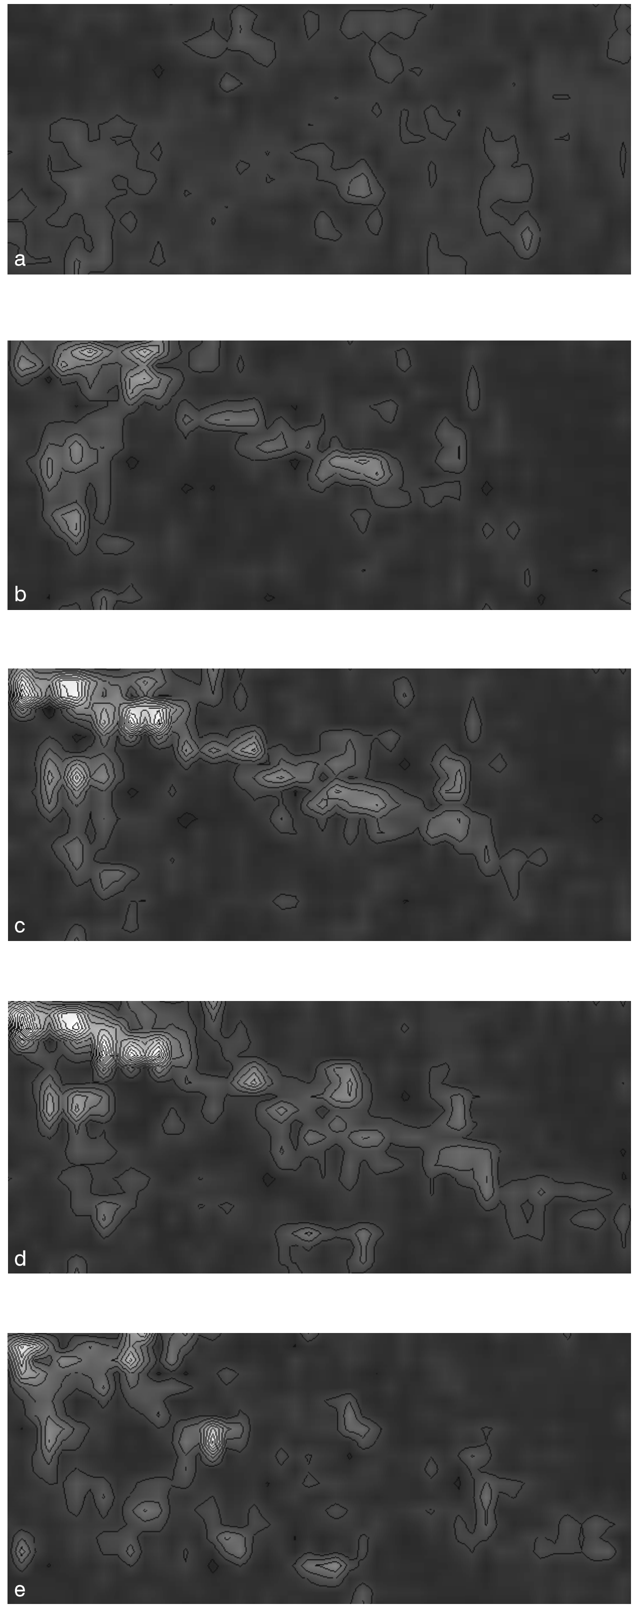

Figure 10. Five consecutive SSE slices through Worcester site.

references

Blake, V.S. 1995 Image processing and interpretation of ground penetrating radar data. In: J. Huggett/ N. Ryan (eds), *Computer Applications and Quantitative Methods in Archaeology 1994*, 175-180, BAR International Series 600, Oxford: Tempus Reparatum.

Daniels, D.J., D.J. Gurron, H.F. Scott 1988 Introduction to subsurface radar, *IEE Proceedings* 135 (4), 278-320.

Fletcher, M., R. Spicer 1990 Ridge and furrow – regression and fourier. In: P. Reilly/S.Rahtz (eds), *Communication in archaeology: a global view of the impact of information technology. vol. 3: late papers*, 31-44, The World Archaeological Congress.

Milligan, R., M. Aitken 1993 The use of ground-probing radar within a digital environment on archaeological sites. In: J. Andresen/T. Madsen/I. Scollar (eds), *Computing The Past, Computer Applications and Quantitative Methods in Archaeology CAA92*, 21-32, Aarhus: Aarhus University Press.

Spicer, R.D. 1991 *Applications of graphical techniques in archaeology and their implications on methodoly*. PhD Thesis, School of Computing, Staffordshire University, Beaconside.

Stratascan 1994 A report for the University of Staffordshire on a geophysical survey carried out at Stapeley Hill.

Jon Bradley
Staffordshire University
School of Computing
Beaconside
Stafford
United Kingdom
e-mail: J.C.Bradley@soc.staffs.ac.uk

Mike Fletcher
Staffordshire University
School of Computing
Leek Road
Stoke
United Kingdom
e-mail: M.Fletcher@soc.staffs.ac.uk