

11

Algorithms for the enhancement and reconstruction of archaeological data sets

J. E. Szymanski*

11.1 Introduction

Over the last two decades a range of remote surveying techniques have been gaining steadily in importance as basic tools of archaeological site analysis. The basic theories behind most of these techniques, in particular magnetic and resistivity surveying, were established long ago (Aitken 1963, Clark 1963, Scollar 1969a, Linington 1973) but only limited use has been made of the full theory of such approaches. This is due mainly to the considerable corruption and distortion of the experimental data which results from the practical and physical limitations of the survey process: there is little point in pursuing detailed mathematical simulations when it will be impossible to test the qualities of such simulations against field data.

In recent years technological advances have changed not only the accuracy of the surveying equipment, but also the way in which the experimental data can be recorded. The use of portable computers has begun to replace the error-prone and laborious task of keeping hand-written records, allowing large quantities of field data to be gathered accurately and rapidly. This has led, in turn, to the possibility of subjecting data sets to mathematical analysis much more readily than in the past. Even two decades ago Linington (1968, 1969, 1970a, 1970b, 1971) and Scollar (Scollar & Krückeberg 1966, Scollar 1969b, Scollar 1970a, Scollar 1970b) made heavy use of computer filtering of archaeological site data, but found that the effort of entering the survey data into the computer was considerable. The ease and speed inherent in the use of digital data recorders has greatly eased the investigation of a range of possible algorithms for the automatic processing of site data.

The speed, power and storage capabilities of battery-operated portable computers are increasing steadily. This, together with the ready availability of small and cheap internal hard disks for portable computers, means that there is now little difficulty in *simultaneously* storing data sets from resistivity or magnetic surveys, or possibly also results from electromagnetic surveys, terrain analysis or even ground-penetrating radar.

This paper discusses some ways in which it is hoped that the clarity and interpretability of such pre-excitation survey data sets can be improved considerably. This work forms part of a joint project at York which is aimed at moving towards the concept of *total mapping*, whereby a large amount of useful information can be obtained for a given site without excavation and in a non-destructive and non-intrusive fashion.

* Department of Electronics
University of York,
Heslington
York

11.2 'Classical' data processing

Whatever survey technique is employed, the data are degraded by non-linearities due to anisotropies in the surface soil, by convolution with an anisotropic and object-specific point-spread function, and by large amounts of different types of noise corruption. Most methods provide little quantitative information about even the coarser detail of depth of an archaeological artefact or feature, indicating only that a feature of some broad description may be found in a certain, fairly loosely defined, area. More specifically, the following problems are associated with the current remote sensing techniques: (a) The surveying technique itself can be slow and inaccurate, owing to the demands of making many measurements over a precisely defined grid on uneven terrain. (b) The resolution available with most approaches is very low. (c) Even major features can be lost in the high noise background and a number of false 'ghost' images can be created by the noise itself. (d) The signal variations obtained are ill-defined and unspecific. (e) Any enhancement of the data is usually performed at a later stage using a mainframe computer. This can be a formidable task, involving large quantities of data being transferred from the on-site data logger (still often human!). This leads to overly long elapsed times between the initial survey and the presentation of the processed data. (f) It is difficult for a layman to understand a typical survey 'image', even after processing.

Common *linear* techniques of data manipulation include origin shifting and data-scaling for contrast enhancement; Fourier analysis and high- or low-pass frequency filtering; deconvolution of blurred data and cross-correlation for the extraction of characteristic signals from a background. These methods are efficacious under circumstances where the noise level is low or when the frequency spectrum of the noise/corruption process is well separated from that of the desired signal. In general, however, these approaches will tend to amplify the high-frequency noise present in the measured data and will give a less than satisfactory overall enhancement.

Non-linear approaches such as spike removal, median filtering, and local averaging are widely used and are considerably more robust when used for the enhancement of noisy systems. The major deficiencies of these techniques are their arbitrary nature (there are a wide selection of different types of median filter, for example: see Arce *et al.* 1986); the standard assumption that the data have been measured on a exact regular grid; and the asymmetric treatment of data at the edge of the grid. At a more advanced level, non-linear adaptive histogram equalization is a non-linear contrast enhancement approach which has been successfully used for the digital processing of medical images (Cocklin *et al.* 1983) among others, but which can introduce spurious features and higher noise levels into the processed image (Pizer *et al.* 1987).

Inverse filtering or transfer function modification (Freiden 1975) employs a bandwidth-limited pseudo-inverse filter in Fourier space while Wiener filtering (Pratt 1978) constructs a restoration which optimises a mean-square measure of an inner-product form of error for the image. Both methods require some advance knowledge of the frequency spectrum of both the 'true' image and the corrupting noise processes. The main problem with such approaches is an inherent ill-conditioning which leads to a lack of smoothness in the processed results. They also, however, produce large negative regions in the enhanced image, which are physically impossible if, in reality, the measurements correspond to the intensity of some experimental quantity.

In short, a wide range of well-established image- and data-processing algorithms of varying complexity are in common use: many require assumptions about the form of the desired image and the inherent noise component; all contain variable parameters which can vary wildly from application to application. This makes the automatic application of enhancement algorithms a less-than-straightforward process. Most importantly, however, these techniques tend to cope badly with high

levels of noise, and produce characteristic artefacts such as ripples and negative areas in the enhanced image. Overall, these facts make it difficult to subject a processed image to quantitative analysis: it may be easier to examine a processed image by eye, but it is also easy for information pertaining to fine structure to be suppressed and for completely spurious and misleading features to be introduced.

Two methods are described below which perform rather better at the low signal-to-noise levels encountered typically in archaeological applications. Both methods are still imperfect in that they also require a range of assumptions and parameters, but they have the advantage of offering the potential for extracting interpretable information from archaeological data that would be difficult to enhance using more standard techniques.

11.3 Simulated Annealing

Simulated annealing (SA) (Kirkpatrick *et al.* 1983, Kirkpatrick 1984) is a novel technique of combinatorial optimisation which is based on the simultaneous use of ideas from probability theory, statistical mechanics, thermodynamics and condensed matter physics. In its basic form a data set is viewed as being equivalent to a physical system with a variable structure to which a structure dependent 'temperature' can be assigned. The measured data set is viewed as a single, non-equilibrium configuration within a vast set of possible configurations, known as an ensemble.

In a practical sense, annealing is a process whereby a system in a highly non-equilibrium state (such as the surface of a freshly fractured crystal) is heated rapidly and then 'repaired' (allowed to settle into a low-energy stable configuration) by undergoing a gradual controlled decrease in temperature. In the optimisation context, if it is possible to create a model which allows an effective temperature to be assigned to each possible configuration, then the optimal set can be viewed as that with the lowest temperature. Pictorially, the noise within the measured data can be viewed as thermal noise, or as particles which have 'evaporated' from the desired image: the algorithm then attempts to reverse the (statistical) process of evaporation of a solid and hence freeze out an enhanced image.

Geman and Geman (1984) explain the algorithm in terms of the equivalence of the Gibbs functions of thermal physics and the Markov random fields of probability theory. They put the algorithm on a detailed mathematical footing as a stochastic relaxation approach, whereby random changes are made to the data set with a probability governed by the overall temperature of the image and by the local, environment-dependent, energy state of any pixel within the image. A potential function

$$U(I) = \sum_i^N \sum_j^N V(i, j) + \sum_i^N \frac{(D_i - I_i)^2}{2\sigma^2} \quad (11.1)$$

is set up where N is the number of pixels in the image; D_i is the i th pixel of the measured image; I_i is the i th pixel of the conjectured enhancement; $V(i, j)$ is negative for neighbouring pixels which are alike, positive for neighbouring pixels which are unlike, and zero otherwise; and σ is the standard deviation of the (assumed Gaussian) noise. The image is assigned some starting temperature and random perturbations are applied on a pixel-by-pixel basis. The two terms of the potential function react very differently to any change in a pixel: the first term will encourage an explicit continuity of large areas of the enhanced image and will lower the potential energy function only when a pixel change is in line with this, while the second demands that there is a cost to any such change if it goes against the measured data. The two terms play off against each other: if the first were absent the enhanced image would

be identical to the measured data; if the second were absent the optimal enhancement would be a uniform shade.

The annealing aspect of the algorithm arises with regard to ensuring that the process does not become trapped in a false, sub-optimal, solution due to some artefact of the data or calculation. This is achieved by also allowing changes which increase the energy function, but with a probability that is governed by the effective temperature. In a physical sense this means that large changes in the data can be made initially, while the system is 'hot', ensuring that the algorithm is able to sample a large volume of the multi-dimensional parameter space and avoid being deceived by purely local structure. As the temperature is lowered the system becomes less likely to make any changes which are energetically expensive, but continues to accept all moves with negative contributions. The net effect is slowly to freeze out an image which is more continuous in form: specular structure is far too costly energetically and is almost completely suppressed by the approach. An important aspect of this approach is the way in which the temperature is reduced: too quickly and the result will certainly be sub-optimal; too slowly and the algorithm becomes computationally unwieldy. Geman & Geman 1984 suggest a suitable temperature profile which cools the system rapidly at first, followed by a much longer and slower approach to the enhanced result.

Because of the recursive nature of the algorithm SA is computationally expensive. Each pixel potential has to be calculated for an entire image and the entire process repeated hundreds of times in an iterative fashion. The results can be worth the effort, however: the approach can be applied successfully to quite general multi-level images with low signal-to-noise ratios. Geman and Geman show results for a series of test images which have been subjected to additive and multiplicative noise, blur and non-linear transformations. For illustrative purposes one of their results is reproduced in Fig. 11.1. This shows the effect of the above algorithm when applied to an image (Fig. 11.1a) corrupted by additive noise (Fig. 11.1b) to produce an enhanced image after 1000 iterations (Fig. 11.1c). Fig. 11.1d shows the effect of introducing a modification into the potential which allows for the introduction of explicit energy-dependent *line processes* into the calculation. These provide a way of introducing terms which incorporate simple explicit edge discontinuities into the image at no extra energy cost but which provide an energy penalty for the more complicated forms of edge structure relative to straight edges.

The SA algorithm has been implemented successfully in a parallel formalism by Murray *et al.* 1986, bringing the time for a complete 64×64 calculation down from hours to only seconds. They found quantitative differences in the results due to the different updating procedures involved in the serial and parallel algorithms, but showed that the approach can be used for rapid image enhancement using existing hardware technology.

The advantages of this stochastic algorithm are that it is very effective at suppressing specular noise, can cope with realistic grey-level images and does not produce any of the ripples or negative areas that are characteristic of some linear techniques. The disadvantages are that it is computationally expensive and has a very arbitrary mathematical formulation (for although the basic concept has a physical analogue, the quantitative nature of the potential function is based on a series of less secure assumptions about the nature of the desired image).

11.4 Maximum Entropy

The Maximum Entropy Method (MEM) was first applied to a range of image processing applications just over a decade ago by Gull & Daniell 1978, although its origins go back several decades further. Since then the method has been developed intensively




Figure 11.1: An illustration of image enhancement by simulated annealing, after Geman & Geman 1984. (a) the original image; (b) the image (a) corrupted by additive noise; (c) image (b) enhanced after 1000 iterations; (d) image (b) enhanced using a modified potential favourable to the presence of straight lines.

by a research group at Cambridge (Kemp & Skilling 1982, Gull & Skilling 1983, Skilling 1984, Gull & Skilling 1985, Skilling & Gull 1985, Skilling 1986, Gull & Newton 1986) and has gained a considerable following. The approach considers a chi-squared statistic

$$\chi^2 = \sum_i^N \frac{(D_i - I_i)^2}{\sigma_i^2} \quad (11.2)$$

where N , D_i and I_i are as defined above, and σ_i is the noise variance at the i th pixel. For some preset value χ_o^2 it is possible to build a set of *feasible images*, each of which satisfy $\chi^2 \leq \chi_o^2$. Out of this huge set of images which are 'not too far' away from the original data it is necessary to select one image which is 'best' in some sense. The approach taken within this algorithm is to find the feasible image which possesses the largest *configurational entropy*

$$S = - \sum_i^N p_i \log \left[\frac{p_i}{m_i} \right] \quad (11.3)$$

where p_i is the normalised intensity of the i th pixel and m_i is some measure of *initial (prior) information* in this pixel. There has been some debate as to the appropriate form of the entropy expression but the Shannon formula of (11.3) has emerged as the most general (Gull & Skilling 1985). In effect, it is a measure of the information in the set of pixels $\{p_i\}$ in the presence of some prior model set $\{m_i\}$. The action of maximizing S subject to the image being chosen from the feasible set corresponds to choosing the image which makes least use of any further assumptions or information outside of that contained in the data and the prior information m_i . In this sense the MEM is said to be the maximally non-committal, or minimally informative, reconstruction: the chosen enhanced image should only contain structure for which there is some evidence in the measured data or prior data. Spurious, algorithm-dependent, artefacts should not appear.

MEM is on a far more solid theoretical footing than SA and has a number of distinct advantages over all methods so far mentioned. Like SA, and unlike inverse and Wiener filters, MEM can only produce positive reconstructions. Also like SA, MEM will suppress strongly noise and spurious features. Further, MEM is capable of reconstructing finer detail than SA and does not require any prior assumptions about the nature of images and edges. If, however, it is desired to insert prior information into the algorithm, then MEM allows this in a natural fashion via the prior model m_i . If nothing is previously known about the image, then the prior can merely be a uniform distribution: if facts are known about the expected form of image, perhaps from model data, these facts can be encoded consistently into the algorithm via the prior. Gull & Newton 1986 show the value of prior information when applied to a practical problem of tomographic imaging.

The computational problem amounts to a constrained maximization of the configurational entropy subject to the chi-squared measure of deviation from the experimental results. This again requires a highly iterative approach and is computationally expensive. A numeric coprocessor or transputer is of considerable use in reducing the elapsed time for a reconstruction.

Perhaps the greatest advantage of MEM from an archaeological point of view is that, unlike most other methods, it does not require that a full set of measurements be taken over an evenly-spaced rectangular grid. The approach treats all data points consistently from an information-theoretic point of view. In this sense MEM can be viewed as the approach which, given unevenly-spaced and sparsely-sampled data, will produce a fit to those data which does not treat end points or isolated measurements any differently from the rest: it uses all the information implicit

in all measurements (including their stated positions) and no more. This point is illustrated in Figs. 11.2 and 11.3. These are reproduced from Skilling & Gull 1985. Fig. 11.2a shows an original photograph of 'Susie' and Fig. 11.2b, d show this image when degraded by convolution with a point spread function of radius six pixels and respectively subjected to heavy and light additive noise. The corresponding MEM reconstructions are shown in Figs. 11.2c and 11.2e. Fig. 11.3 then looks at the effect of sparse sampling on the MEM reconstruction of Fig. 11.2d. Figs. 11.3a,c,e show the degraded image 11.2d further corrupted by simply randomly discarding 50%, 95% and 99% of the data respectively. Fig. 11.3b,d,f show the corresponding MEM reconstructions. It is, at first sight, amazing that a recognizable reconstruction can be obtained from only a few percent of noisy, blurred data.

11.5 Overview

Both SA and MEM are stochastic algorithms with roots in Bayesian statistics that operate on experimental data sets in a non-linear iterative fashion. Both algorithms offer significant advantages over linear and median filtering approaches with regard to noise suppression and positivity constraints. The interpretability of the enhanced images is also usually considerably improved.

The main cost of both algorithms is processing power, but from an archaeological point of view the necessary expenditure may be well worthwhile. MEM, in particular, offers two extremely important advantages: (i) the ability to insert into the algorithm any model data via prior information; and (ii) the ability to cope with sparse and unevenly sampled noisy data.

Feature (i) has potential for the consistent treatment of data arising from simulation models or multiple surveys. Bartlett 1987 has given a clear exposition of the need for increased work on the use of simulation models for archaeological surveys. Imai *et al.* 1987 have employed simultaneous resistivity and ground-probing radar surveys to show the benefits of a comparative qualitative study of the two sets of data. With MEM it should be possible to go further and to use the results of one survey as the prior model for the enhancement of the results of a second: no other method offers the possibility of achieving this in any quantitative fashion.

Given the usual circumstances of an archaeological site survey, feature (ii) of MEM has considerable potential in that it need not matter, say, if a modern road has been laid across the survey site: the algorithm can consistently combine survey data from either side of the obstacle to construct a reconstruction over the *entire* area. Neither is it necessary to produce a regular grid of measurements if this is inconvenient: as long as some accurate method of location (microwave position measurement, for example) is employed then the measurements can be located as desired. An important aspect of this is that it is possible to take an increased number of measurements in areas suspected to be of particular interest, thus increasing the local information density and the reliability of the enhanced results for that area.

A joint project has been initiated at York, one aspect of which is to investigate the use of modern stochastic relaxation algorithms in an archaeological context. Early indications are that both the SA and MEM algorithms seem to have wide applicability in archaeological data enhancement. Particularly interesting areas for investigation are the use of archaeology-specific model simulations or multiple surveys to provide prior information for MEM. It is hoped in the future to develop a range of 'cut-down' stochastic algorithms, suitable for immediate preliminary on-site analysis using commonly available portable computers.




Figure 11.2: An illustration of image enhancement by the Maximum Entropy Method (MEM), after Skilling & Gull 1985. (a) the original image; (b) image (a) degraded by convolution and subjected to heavy additive noise; (c) image (b) reconstructed by MEM; (d) image (a) degraded by convolution and subjected to light additive noise; (e) image (d) reconstructed by MEM;




Figure 11.3: An illustration of the effect of sparse sampling on a MEM reconstruction, after Skilling & Gull 1985. (a), (c), (e): image 1.2(d) further corrupted by randomly discarding 50%, 95% and 99% of the data respectively; (b), (d), (f): the corresponding MEM reconstructions.

Bibliography

- AITKEN, M. 1963. "Magnetic Location". in Brothwell, D., Higgs, E., & Clark, G., (eds.), *Science in Archaeology: A Survey of Progress and Research*, pp. 681–694. Thames and Hudson.
- ARCE, G. R., N. C. GALLAGHER, & T. A. NODES 1986. "Median filters: Theory for one- and two-dimensional filters", *Advances in Computer Vision and Image Processing*, 2: 89–166.
- BARTLETT, A. 1987. "Computing and archaeological geophysics". in Richards, J. D., (ed.), *Computer usage in British Archaeology*, Occasional Paper 1, pp. 55–61. Institute of Field Archaeologists, Birmingham.
- CLARK, A. 1963. "Resistivity Surveying". in Brothwell, D., Higgs, E., & Clark, G., (eds.), *Science in Archaeology: A Survey of Progress and Research*, pp. 695–707. Thames and Hudson.
- COCKLIN, M. L., A. R. GOULAY, P. H. JACKSON, G. KAYE, I. H. KERR, & P. LAMS 1983. "Digital processing of chest radiographs," *Image and Vision Computing*, 1 (2): 67–78.
- FREIDEN, B. R. 1975. "Image Enhancement and Restoration". in Huang, T. S., (ed.), *Topics in Applied Physics Volume 6: Picture Processing and Digital Filtering*, pp. 177–248. Springer-Verlag.
- GEMAN, S. & D. GEMAN 1984. "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6): 721–741.
- GULL, S. F. & G. J. DANIELL 1978. "Image reconstruction from incomplete and noisy data", *Nature*, 272: 686–690.
- GULL, S. F. & T. J. NEWTON 1986. "Maximum entropy tomography", *Applied Optics*, 25 (1): 156–160.
- GULL, S. F. & J. SKILLING 1983. "The Maximum Entropy Method". in Roberts, J. A., (ed.), *INDIRECT IMAGING: Measurement and processing for indirect imaging (Proceedings of an International Symposium, Sydney, Australia, 30 August–2 September 1983)*, pp. 267–279.
- GULL, S. F. & J. SKILLING 1985. "The Entropy of an Image", *Maximum-Entropy and Bayesian Methods in Inverse Problems*, pp. 287–301.
- IMAI, T., T. SAKAYAMA, & T. KANEMORI 1987. "Use of ground-probing radar and resistivity surveys for archaeological investigations", *Geophysics*, 52 (2): 137–150.
- KEMP, M. C. & J. SKILLING 1982. "Image reconstruction by Maximum entropy", in *Proceedings of the International Conference on ELECTRONIC IMAGE PROCESSING, 26-28 July 1982, University of York, UK, (IEE, Conference publication number 214)*, pp. 174–178.
- KIRKPATRICK, S. 1984. "Optimisation by Simulated Annealing: Quantitative Studies", *Journal of Statistical Physics*, 34 (5/6): 975–986.
- KIRKPATRICK, S., C. D. GELATT, & M. P. VECCHI 1983. "Optimisation by Simulated Annealing", *Science*, 220 (4598): 671–680.

- LININGTON, R. E. 1968. "The Rome computer system for treating archaeological survey results. First part", *Prospezioni Archeologiche*, 3: 19-36.
- LININGTON, R. E. 1969. "The Rome computer system for treating archaeological survey results. Second part", *Prospezioni Archeologiche*, 4: 9-58.
- LININGTON, R. E. 1970a. "A first use of linear filtering techniques on archeological prospecting results", *Prospezioni Archeologiche*, 5: 43-54.
- LININGTON, R. E. 1970b. "Techniques used in archaeological field surveys", *Phil. Trans. Roy. Soc. Lond. A*, 269: 89-108.
- LININGTON, R. E. 1971. "Further tests on non-symmetrical filtering systems", *Prospezioni Archeologiche*, 6: 9-20.
- LININGTON, R. E. 1973. "A summary of simple theory applicable to magnetic prospecting in archaeology", *Prospezioni Archeologiche*, 7/8: 9-60.
- MURRAY, D. W., A. KASHKO, & H. BUXTON 1986. "A parallel approach to the picture restoration algorithm of Geman and Geman on an SIMD machine", *Image and Vision Computing*, 4 (3): 133-142.
- PIZER, S. M., E. P AMBURN, J. D AUSTIN, R. CROMARTIE, A. GESELOWITZ, T. GREER, B. T. H. ROMENY, J. B. ZIMMERMAN, & K. ZUIDERVELD 1987. "Adaptive Histogram Equilization and its Variations", *Computer Vision, Graphics and Image Processing*, 39: 355-368.
- PRATT, W. K. 1978. *Digital Image Processing*. Wiley, New York.
- SCOLLAR, I. 1969a. "A program for the simulation of magnetic anomalies of archaeological origin in a computer", *Prospezioni Archeologiche*, 4: 59-85.
- SCOLLAR, I. 1969b. "Some techniques for the evaluation of archaeological magnetometer surveys," *World Archaeology*, 1: 77-89.
- SCOLLAR, I. 1970a. "Fourier Transform Methods for the Evaluation of Magnetic Maps," *Prospezioni Archeologiche*, 5: 9-41.
- SCOLLAR, I. 1970b. "Magnetic methods of archaeological prospecting—advances in instrumentation and evaluation techniques", *Phil. Trans. Roy. Soc. Lond. A*, 269: 109-119.
- SCOLLAR, I. & F. KRÜCKEBERG 1966. "Computer treatment of magnetic measurements from archaeological sites", *Archaeometry*, 9: 61-71.
- SKILLING, J. 1984. "The maximum entropy method", *Nature*, 309: 748-749.
- SKILLING, J. 1986. "Theory of Maximum Entropy Image Reconstruction". in Justice, J. H., (ed.), *in Maximum Entropy and Bayesian Methods in Applied Statistics, (Proceedings of the Fourth Maximum Entropy Workshop, University of Calgary, 1984)*, pp. 178-193. CUP, Cambridge.
- SKILLING, J. & S. F. GULL 1985. "Algorithms and Applications". in Smith, C. R. & Grandy, W. T., (eds.), *Maximum-Entropy and Bayesian Methods in Inverse Problems*, pp. 83-132. Reidel.