

32 The provenancing of archaeological ceramics: a Bayesian approach

C. E. Buck

32.1 INTRODUCTION

The investigation of artefacts with a view to determining their provenance is now a highly sophisticated and complex area of study in archaeology. There are many reasons why archaeologists wish to investigate the original provenance of finds on a site or in a geographical region. These reasons usually relate to the desire to understand the interrelationships of early societies. Similar objects may get to different locations by a wide range of mechanisms; the nature of which may reveal a great deal about society at the time. For example, two similar objects may be found on different sites because they were both made at the same location and then traded, or because itinerant workers travelled to different sites making similar objects at each site, or because one is a copy of the other made by a craftsman who admired someone else's work but who had no cultural link to the other worker at all. Thus, for example, a controversy has developed as to whether Pictorial Pottery was actually made on Cyprus by Mycenaean migrants or was imported from mainland Greece. Catling *et al.* (1963), Catling *et al.* (1978) and Jones (1986:544–548) have argued that, on the basis of chemical compositional and archaeological evidence, the latter was the case.

The varied nature of the mechanisms that lead to similar items appearing at different locations has resulted in the adoption of several different methods for investigating them. Such types of investigation include: typological and stylistic investigation of the whole object, technological investigation of the whole object and/or thin-sections or small samples and chemical compositional analyses of the material from which the objects are made. The basic premise is that if similar objects were originally manufactured at

the same site, during the same time period, by the same techniques, and using the same raw materials, then we would expect to find similarities in style, typology and chemical composition. If however the objects were manufactured during the same time period using the same techniques, but at different sites using different raw materials, then they may be visually similar, but have different chemical compositions. In situations where it is difficult to assign provenance on stylistic or typological grounds, we would hope to gain greater insight through chemical compositional analyses. Thus, for example, when the archaeologists at Maa, Cyprus wished to investigate whether or not storage jars of a particular type were of local manufacture or were imported, they submitted samples for chemical compositional analyses. Jones and Vaughan (1988) have argued, on the basis of their chemical compositional analyses, thin-section investigations and on the report by the pottery expert, that the jars found at Maa were not manufactured locally, but came from elsewhere in Cyprus or even perhaps from further afield. Chemical compositional data are most commonly obtained for the purposes of provenancing ceramics and it is to these that we will refer specifically in the remainder of this paper; it should be noted, however, that the methodology may readily be applied to other types of objects.

32.2 AN OUTLINE OF THE PROBLEM

There are three main types of information commonly collected to aid in the investigation of ceramic provenance:

- 1) archaeological information (for example context, shape, decorative style, date, etc);

- 2) technological information (for example slip, surface treatment, fabric, manufacturing methods, etc);
- 3) chemical composition of the fabric of the pottery.

The results of investigations of an archaeological and/or technological nature usually take the form of written reports by the subject matter experts with a range of conclusions explicitly stated. Such conclusions might take the form of statements about the archaeological forms represented, the date ranges for the various forms, which of the pottery fragments (sherds) are from vessels likely to have been manufactured using the same methods, which have similar methods of decoration, which have similar (or different) geological inclusions in their fabric and which are from vessels that have been fired under similar (or different) conditions.

The results from chemical compositional analyses usually take the form of large amounts of quantitative data (in percent or parts per million) for a wide range of elements. Each sample analysed may result in as many as 30 variables, one for each element examined. Within data of this sort we would expect different provenances to show up in the form of groups of sherds with similar chemical compositions, we therefore need a statistical method that will allow us to recognise evidence of this. There are now a number of texts which offer recommendations about selecting such methods, eg. Bieber *et al.* (1976), Wilson (1978) and Pollard (1982, 1986). Traditionally such methods have not incorporated other types of information with the chemical variables, but have analysed the data "blind".

We should make it clear at this point that the statistical analysis of chemical compositional data takes two different routes, one in situations where samples from likely known sources are available and one where they are not. The former is the more straight forward statistically and the recommended approach is usually to use *discriminant analysis* (Pollard 1986:65–66). Fillieres *et al.* (1983) used this method to investigate figurines, pottery and workshop material from the Athenian Agora. The more statistically complex problem arises when we have no known source material with which to compare our current samples and in these circumstances the most widely recommended approach is that of *cluster analysis*.

The term cluster analysis does not apply to just one statistical technique, but to a wide range of different techniques with the same basic aim. All such techniques analyse multivariate data and, on

the basis of some type of distance measure, produce subdivisions. Most of the works, cited above, which give recommendations about selection of clustering methods conclude that standard methods available within commercial statistical packages have wide applicability to archaeological provenancing problems. However, some workers, perceiving weaknesses in these standard approaches, have investigated other methods. Rice & Saffer (1982) and Vitali (1989) considered alternative approaches to the analysis of chemical compositional data for pottery provenancing and Baxter (1991) considered the problem with a view to investigating data from mediaeval glass. The perceived weaknesses referred to above are varied and often depend on the type of archaeological objects being investigated. One common theme is that the use of traditional statistical techniques has resulted in much information being ignored or used only for the purposes of comparison after the chemical compositional data has been analysed (Rice & Saffer 1982:398). It would be extremely unusual when undertaking chemical compositional analyses for the purposes of provenance studies not to have information available, prior to analysing the data, about suspected cluster membership. As indicated above, such information commonly arises from archaeological or technological studies which are almost invariably carried out prior to the chemical analyses.

In this paper we propose a Bayesian approach to the investigation of such problems. We present a method for dividing the data into homogeneous groups based on discriminant analysis but which incorporates the archaeological and or technological information from the subject matter experts. We believe that the resultant method gives a coherent approach to the combination of data and prior information from a wide variety of sources.

32.3 THE NATURE OF THE SUBJECT MATTER EXPERTS' INFORMATION

The type of investigations that the ceramics have undergone prior to chemical compositional analysis will vary depending upon the scale of the study. We could reasonably expect, however, that all sherds submitted for chemical analyses will, as a minimum, have been studied first by a pottery expert. This will mean that a pottery report exists. The nature of the information provided in such a report varies quite considerably, but will usually include a detailed description of each "type" of pottery represented in the collection. By "type" we mean archaeological "type" which is usually

defined on stylistic grounds. In some cases there will, in addition to this report, be information resulting from thin-section microscopic analyses (cf. Day 1988). Such analyses can allow closer definition of types than is possible from the whole sherds alone and may allow further subdivisions of these types. Identification and quantitative description of petrographic inclusions, in many cases only possible in thin-section under a microscope, are extremely valuable in the provenancing of pottery. For example, it is clear from Jones and Vaughan (1988), mentioned above, that the study of thin-sections formed an integral part of the investigation of the jars and indeed aided greatly in the the general conclusions that were made about their provenance.

As a result of such detailed investigation by subject matter experts, it should be possible to group together some samples likely (but not certainly) to have the same provenance. Conversely, it should be possible to identify some samples which have very different archaeological, technological or petrographic types and so are unlikely to be from the same provenance (although, of course there may be a small chance that they could be).

We should emphasise, however, that the evidence provided by the pottery expert, the thin-section analyst, and the chemist do not necessarily point to exactly the same groupings in the data. We must be very careful when defining various groups based on the experts' knowledge to establish that they really are likely to be different on compositional grounds and not just stylistic ones. For example, if several different types of vessel produced at the same site appear in thin-section to have the same matrix, then this would suggest that the samples probably ought to be considered as being part of the same group for the purposes of provenance studies. Conversely, samples from vessels of the same archaeological or technological type may well, in thin-section, have different matrices. Such sherds probably ought to be allocated to different groups if they are to aid in the interpretation of the chemical compositional data for provenancing purposes.

32.4 THE BAYESIAN APPROACH

In order to make use of the Bayesian paradigm, we must have a statistical model which involves a vector, θ , of unknown parameters and the sample data, x . The model defines a relationship between θ and x which gives rise to a *likelihood function*, $l(x;\theta)$. The unknown parameter vector,

θ , is a realisation of a random variable Θ which has a prior density, $p(\theta)$. Inferences about θ are then based upon knowledge of the posterior density, $p(\theta|x)$, which is evaluated using Bayes theorem

$$p(\theta|x) = \frac{l(x;\theta)p(\theta)}{\int l(x;\theta)p(\theta)d\theta} ,$$

where the integration is carried out over the appropriate range of θ .

In any practical application of this approach, it is clear that we may encounter problems in carrying out the necessary integrations. In fact in very many real situations it proves to be impossible to obtain analytic solutions. To overcome such problems, a range of sampling-based approaches has been developed. Such methods, for example the Gibbs sampler (Geman & Geman 1984 and Gelfand & Smith 1990), are relatively easy to implement and are now gaining in popularity. Many of the early applications of these sampling-based approaches were in the field of image processing, for example Geman & Geman (1984) and Besag (1986). From a Bayesian perspective there are many similarities between image processing and cluster analysis and readers are referred to Besag (1986) for an introduction to the original methodological ideas. For archaeologists wishing to gain some insight into the use of Gibbs sampling, Buck & Litton (1991) outlined its use in investigating two types of archaeological problem; one involving the clustering of clay pipe data and the other the seriation of artefactual data.

32.4.1 The statistical model

A Bayesian clustering method for use with archaeological data has been suggested before (Buck & Litton 1991). Here, however, we suggest a rather different formulation of the model which allows for the experts' knowledge to be included.

We now define terminology, to make a clear distinction between **groups** and **clusters** for the remainder of this paper. Subdivisions of the data made by the subject matter experts will now be described as **groups**. Subdivisions made using the Bayesian approach and hence being formed using both the experts' knowledge and the chemical compositional data will be described as **clusters**.

Firstly let us suppose that, on the basis of the experts' investigations of the sherds in a particular assemblage, it is possible to divide the n sherds into J initial groups where group j ($j = 1, 2, \dots, J$) has n_j members. It is quite possible that some groups will have only one member. This al-

lows for sherds which the experts are unable to group with any others due perhaps to lack of knowledge about them, or to evidence that they are unlike any of the other sherds in the assemblage. In this way we can then represent the experts' knowledge in terms of a vector of integers, δ , where $\delta_i = j$ if the experts initially assign object i to group j . Secondly suppose that the maximum number of possible provenances is L and that for sherds from the l th provenance ($l = 1, 2, \dots, L$) the vector, x , representing the chemical compositional data has probability density function $f_l(x)$. We assume that, conditional on the provenance, the chemical compositions of the sherds are independent. That is, for two sherds i and j with chemical composition vectors x_i and x_j , x_i and x_j are independent of one another. Furthermore we assume that the chemical compositional information and subject matter experts' information are independent. That is, we assume that the subject matter experts assign their values to δ without first having access to the chemical compositional data.

Suppose that the $f_l(\cdot)$ are known and that no information is available from the subject matter experts, then using classical discriminant analysis (see Mardia *et al.* 1982) we would assign sherd i to provenance l if

$$f_l(x) = \max_{k=1,2,\dots,L} f_k(x).$$

If prior information were available regarding the assignment, then we would use a Bayesian discriminant rule which allocates sherd i to provenance l if

$$\pi_l f_l(x) = \max_{k=1,2,\dots,L} \pi_k f_k(x)$$

where π_l is the prior probability of assigning a sherd to provenance l . However, in provenancing, we have no information regarding the π_l and so in our method we are going to use the subject matter experts' opinions to provide information about them. There are many ways of doing so one of which is as follows.

Consider sherd i which has been assigned by the subject matter expert to group j . Suppose that all the other sherds, except sherd i , have been assigned by some allocation rule to one of the L provenances. Then we allocate sherd i to provenance l with probability proportional to $\pi_{il}(\delta) f_l(x_i)$ where $\pi_{il}(\delta)$ reflects how many of the sherds originally assigned by the subject matter experts to the same group as sherd i are currently allocated to provenance l . To reflect this, $\pi_{il}(\delta)$ is proportional to $e^{\gamma n_{il}}$, where γ is a constant which

needs specifying. γ is effectively a weighting factor for the subject matter experts' information. If the experts are certain that their information is of high quality, then γ will be large and conversely, if they are unsure about their ability to assign the sherds to initial groups, then the value of γ is small.

32.4.2 Implementation using the Gibbs sampler

The Gibbs sampler is a method of stochastic simulation that has revolutionised Bayesian inference by allowing complex models to be analysed readily. There are now numerous papers on the theory and application of the technique. We refer the reader to Gelfand & Smith (1990) and to the other papers cited above.

For the purposes of this paper, we will assume that a transformation is available that will allow us to transform the raw, multi-element, chemical compositional data to normality (a number of authors have addressed this problem and several such transformations have been applied, for a summary of the arguments see Pollard 1986). Suppose that sherd i comes from provenance l ($l = 1, 2, \dots, L$) then its vector of (transformed) chemical compositions, x_i , is normal with mean μ_l and common covariance matrix Σ . (The assumption of common covariance between provenances has not been widely tested by archaeologists, pottery experts, chemists or statisticians and, should it prove not to be reasonable, will require reformulation of this aspect of the model.)

We assume that the prior density for μ_l has a normal distribution with expectation η_l and covariance matrix $\rho_l^{-1}\Sigma$. Assume also that Σ^{-1} , the inverse of the covariance matrix, has a Wishart distribution with β degrees of freedom and precision matrix τ . Now, let $\phi_i = l$ if the i th sherd is allocated to cluster l . Then the conditional densities are

(i) $\Sigma^{-1} | X, \delta, \gamma, \phi, \mu_1, \mu_2, \dots, \mu_L \sim$ Wishart with degrees of freedom $\beta + n_1 + n_2 + \dots + n_L$ and precision matrix

$$\tau^* = \tau + \sum_{i=1}^L s_i + \sum_{i=1}^L \frac{\rho_i n_i}{\rho_i + n_i} (\eta_i - \bar{x}_i)(\eta_i - \bar{x}_i)'$$

(ii)

$$\mu_i | X, \delta, \gamma, \Sigma, \phi, \mu_i (i \neq l) \sim N\left(\frac{\rho_i \eta_i + n_i \bar{x}_i}{\rho_i + n_i}, (\rho_i + n_i)^{-1} \Sigma\right)$$

and

(iii)

$$p(\phi_i | X, \delta, \mu_1, \mu_2, \dots, \mu_L, \Sigma, \gamma) =$$

$$p_{ii} \propto \pi_{ii}(\delta) \exp\left[-\frac{1}{2}(x_i - \mu_i)' \Sigma^{-1}(x_i - \mu_i)\right]$$

where X is the data matrix whose rows are composed of the x_i vectors and n_l , \bar{x}_l and $n_l^{-1}s_l$ are respectively the number of sherds in cluster l , the sample mean for cluster l and the covariance matrix for cluster l .

32.5 AN EXAMPLE APPLICATION

We will now illustrate the use of the methodology described above to investigate data from known kiln sites. The data available to us are from tin glazed wares produced at three medieval kiln sites two in Spain at Valencia and Seville and one in Italy at Castelli. The pottery is characterised by being highly decorated, some in metallic copper lustre and some in a wide range of colours (blues, greens, yellows and purples) on a white background. It has long been admired and discussed (Caiger-Smith 1973) and is of particular interest to archaeologists and pottery experts as it is known to have been exported from its sources to other European countries and to America (Hurst *et al.* 1986).

We know that all the samples submitted for chemical analysis are from one or other of the three sites and we know which samples are assigned to which site. There are obviously many different studies that could be made of this data and we do not propose to offer here a comprehensive consideration of all the possibilities. What is offered is an outline of the Bayesian methodology and description of some of the effects of changes in the nature and quality of the subject matter experts' information. It should be noted that, although the data are actual chemical compositional values from real samples, the examples of expert information that are used are all fictitious since the clusters present in the data are in fact known.

Available to us are neutron activation analysis results from 150 samples; 63 from Valencia, 63 from Seville and 24 from Castelli. For each sample we have 15 variables, one for each trace element measured. After imputing the very few missing values, taking natural logarithms results in the data being transformed approximately to normality (Pollard, 1986:68).

The methodology outlined above requires prior information regarding likely cluster means and covariances. In what follows we will use the




Figure 32.1: profile plot of the log data from Valencia.




Figure 32.2: profile plot of the log data from Seville.

Castelli data to provide this information (see below for details) and the statistical analysis described will then only be of the data from Valencia and Seville. For convenience we label the sherds to be clustered 1–126, the Valencia samples being 1–63 and the Seville samples 64–126.

Figure 32.1 and Figure 32.2 show plots of the log data for Valencia and Seville respectively. Each of the samples is indicated by its own number, so that these plots allow us to observe which of the samples are possible outliers. We observe that only two samples have extreme values in more than one element, these are 41 and 42. Some other samples (including 21 and 104) have extreme values in only one element each.

We investigate the data still further by plotting the first two principal components of the standardised log data from Valencia and Seville (Figure 32.3). As we might expect, samples 41 and 42 appear as outliers. We note that by using principal components we are combining the information from all 15 elements and that those samples which only have one extreme elemental concentration do not necessarily appear as outliers in Figure 32.3. Furthermore, it is clear from Figure




Figure 32.3: first two principal components of the standardised log data from Valencia and Seville.

32.3 that on the basis of the data alone (ie without making use of the subject matter experts' knowledge) it is extremely difficult to resolve the two clusters. Indeed several of the sherds from Valencia would certainly be misclassified as coming from Seville.

Turning now to consider the Bayesian analysis of these data, we suggest some possible scenarios for information provided by the subject matter experts.

Scenario 1: of the 126 samples the subject matter experts (correctly) assign 30 to one group and 30 to another group but are unsure about the remaining 66 samples. That is, on the basis of the subject matter experts' judgement two groups can be identified within the data leaving half the sherds ungrouped. Expressing this in terms of the resulting δ vector: $\delta_1 = \dots = \delta_{30} = 1$; $\delta_{31} = \dots = \delta_{63} = 0$; $\delta_{64} = \dots = \delta_{93} = 2$ and $\delta_{94} = \dots = \delta_{126} = 0$.

Scenario 2: the subject matter experts believe that there are four groups present in the data and they subdivide the Valencia group into two groups one of 30 sherds and one of 33 sherds and they do the same with the Seville group. This results in the suggestion of a larger number of subdivisions in the data than are actually present as might be the case if one kiln site is producing pottery of two different styles at the same time. Expressing this in terms of the resulting δ vector: $\delta_1 = \dots = \delta_{30} = 1$; $\delta_{31} = \dots = \delta_{63} = 2$; $\delta_{64} = \dots = \delta_{93} = 3$ and $\delta_{94} = \dots = \delta_{126} = 4$.

Scenario 3: the subject matter experts identify only two groups within the data, but (perhaps

due to similarities in decorative style between the two kilns) in fact believe that half of the sherds from Valencia should be grouped with half of those from Seville and vice versa. This results in two groups being identified one containing 30 sherds from Valencia and 33 from Seville and the other containing 33 sherds from Valencia and 30 sherds from Seville. Expressing this in terms of the resulting δ vector: $\delta_1 = \dots = \delta_{30} = 1$; $\delta_{31} = \dots = \delta_{63} = 2$; $\delta_{64} = \dots = \delta_{96} = 1$ and $\delta_{97} = \dots = \delta_{126} = 2$.

Scenario 4: is an unrealistic scenario included to allow a demonstration of a wide range of types of subject matter expert information. This is the situation where the experts correctly assign all the sherds from Valencia to one group and all the sherds from Seville to another group. This results in two equally sized groups with 63 sherds in each. Expressing this in terms of the resulting δ vector: $\delta_1 = \dots = \delta_{63} = 1$ and $\delta_{64} = \dots = \delta_{126} = 2$.

With the data from the three known kilns and the information contained in the δ vectors for the four different scenarios, we are now in a position to investigate the use of the Gibbs sampler methodology outlined above. It will be clear that there are a number of variables which can have a marked effect on the likely outcome of the simulations. For the purposes of this paper we consider the Bayesian analysis of the data and in all cases fix the values of μ , Σ , γ , β and ρ so that we are reporting only the effects of varying the values in the δ vector. In the case of the results reported here, we obtained the values for μ by taking the sample mean of the Castelli data and adding a value simulated from a uniform distribution in the range minus twice the standard deviation of the Castelli data to plus twice the standard deviation of the Castelli data. This is not entirely satisfactory and so ρ is correspondingly set to 1.0, indicating that we are only able to provide a very approximate estimate for μ . We set Σ equal to the sample covariance matrix of the Castelli data which, since common covariance matrices are assumed in the model, should be a reasonable estimate. We set $\gamma = 1.0$, corresponding to a fairly firm belief in the subject matter experts' information and, to reflect our belief in how accurate Σ is, we set $\beta = 60$.

With these parameter values fixed in this way we now report the effects of investigating the Valencia and Seville data in conjunction with the four different subject matter expert scenarios. We find that both the quantity and the accuracy of the subject matter experts' information have a marked effect on the outcome of the Bayesian analysis. When using scenario 4 convergence is rapid, leading to two main clusters, correctly as-

signing the Valencia and Seville sherds to separate clusters. In addition samples 41 and 42 are commonly assigned to separate clusters possibly indicating that these samples are outliers. When using scenario 1, however, convergence is much slower. The result is still two main clusters, with samples 41 and/or 42 identified as outliers. There is considerably less subject matter expert information available from scenario 4 than from scenario 1 and thus the slower convergence rate when using scenario 1 is entirely to be expected.

In terms of both quantity and accuracy of information, scenarios 2 and 3 are comparable to one another. This is reflected in the way that they influence the outcome of the Bayesian data analysis. In both cases we find that convergence is reached much more slowly than with either of the other two scenarios and that there is some uncertainty about when convergence has occurred. Although the final result is two main clusters with some outliers (centred around samples 41 and 42) it is not uncommon for convergence to appear to have been reached earlier. Without using a convergence measure and without knowing the "true" subdivisions present it would be easy to conclude that there are in fact three main clusters in the data. The "extra" cluster is formed from the Valencia data and is made up of samples from the top right hand corner of the principal components plot (Figure 32.3). Inspection of Figures 32.1, 32.2 and 32.3 indicates that the Valencia data are in fact rather more disparate than those from Seville.

The more disparate nature of the Valencia data could be a reflection of one or more of several different features of the manufacturing process of the ceramics in question. As part of the manufacture of pottery vessels the raw clay is almost always preprocessed in one or more ways. Such preprocessing may not be entirely consistent from one batch of clay to the next. For example, the amount of temper added is usually of an approximately fixed proportion to the clay, but may vary on a batch to batch basis. The magnitude of such variability will depend, amongst other things, on the method used to measure the various quantities of materials required. We do not have access to detailed records about the likely mode of manufacture of the ceramics from which the Valencia samples were taken and so are unable to state whether there is any archaeological interpretation for the "extra" cluster. In any real analysis of this sort, however, dialogue between the statistician and the archaeologist would be necessary in order to allow full assessment and informed interpretation of the statistical results. Ideally, of




Figure 32.4: fusion plot of the results from Ward's method cluster analysis of the log data from Valencia and Seville.

course, if the archaeologist (or pottery expert) has reason to believe in the "extra" group prior to obtaining the data then the information should be included in the δ vector and thus incorporated into the statistical analysis.

We now turn to look at how the Gibbs sampler results compared with those obtained using traditional cluster analysis. We ran the logged Valencia and Seville data through the CLUSTAN package using Ward's method. Looking at the fusion points for this analysis, we find that the Seville cluster begins to form first and that the outliers (samples 41 and 42) in the Valencia data are amongst the last to cluster. Figure 4 shows a plot of the fusion coefficients for the CLUSTAN run which indicates clearly that there are likely to be either 2 or 3 clusters in the data. At the two cluster level running CLUSTAN using Ward's method gives perfect subdivision of the data into Seville sherds and Valencia sherds. At the three cluster level, however, CLUSTAN divides the Valencia cluster into two approximately equal clusters and at no point in the analysis do we pick up the small "extra" cluster that was observed in the Bayesian analysis.

32.6 CONCLUSIONS

We have developed a Bayesian approach to the statistical analysis of chemical compositional data for the purposes of provenancing archaeological ceramics. The method allows us to combine the information contained within the chemical data

with the expert knowledge of the archaeologists and pottery specialists. Such experts have for sometime been able to make assessments of likely provenance groupings on the basis of typology, style, technology and/or thin-section investigations. It is only with the advent of the Gibbs sampler, however, that this statistical approach has become practical.

The author now intends to test the approach using larger and more complex data sets and would be glad to hear from specialists in any of the many fields in which chemical compositional data is collected for the purposes of provenancing. If Bayesian analysis is to be used successfully in these fields, as in any others, true collaboration is required between statisticians and archaeological experts from all relevant fields. We can no longer expect to work in each field in isolation and only pool our results at the "interpretation" stage. When using a Bayesian approach, interpretable results only arise from analysis based upon clearly stated subject matter expert information leading to the development of relevant statistical models and to the elicitation of prior information about the parameters of the model.

Acknowledgements

The author is indebted to Dr W. G. Cavanagh who presented this paper on her behalf at CAA92 in Denmark. While all errors and omissions remain the authors own, this paper would not have been possible without his generous support. The author is also grateful to Dr M. J. Hughes of the British Museum Research Laboratory for making available the data used in the examples section of this paper and for his constructive comments early on in its preparation. In addition the author acknowledges the advice and support of a number of other colleagues in particular Drs C. D. Litton and D. A. Stephens whose comments on earlier drafts of this paper proved invaluable. This work was completed while the author was employed under a research grant from the Science-based Archaeology Committee of the UK Science and Engineering Research Council grant number GR/G42426.

References

- Baxter, M. J.
1991 Principal component and correspondence analysis of glass compositions: an empirical study. *Archaeometry* 33(1):29–41.
- Besag, J.
1986 On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society Series B* 48 (3):259–302.
- Bieber, A. M., D. W. Brooks, G. Harbottle & E. V. Sayre
1976 Application of multivariate techniques to analytical data on Aegean ceramics. *Archaeometry* 18(1):57–74.
- Buck, C. E. & C. D. Litton
1991 A computational Bayes approach to some common archaeological problems. in Kris Lockyear & Sebastian Rahtz (eds.) *Computer Applications and Quantitative Methods in Archaeology, 1990*, British Archaeological Reports, International Series 565 pp. 93–99.
- Caiger-Smith, A.
1973 *Tin Glaze Pottery in Europe and the Islamic World*. Faber and Faber, London.
- Catling, H. W., R. E. Jones & A. Millett
1978 Composition and provenance problems in some LBA pottery found in Cyprus *Report of the Department of Antiquities, Cyprus*:70–90.
- Catling, H. W., E. E. Richards & A. E. Blin-Stoyle
1963 Correlations between composition and provenance of Mycenaean and Minoan pottery. *Annual of the British School at Athens* 58:94–115.
- Day, P.
1988 Ceramic analysis and pottery systems: the case of Minoan Crete. in R. E. Jones & H. W. Catling (eds.) *New Aspects of Archaeological Science in Greece* pp.39–45.
- Fillieres, D., G. Harbottle & E. V. Sayre
1983 Neutron activation study of figurines, pottery and workshop material from the Athenian Agora, Greece. *Journal of Field Archaeology* 10:55–69.
- Gelfand, A. E. & A. F. M. Smith
1990 Sampling based approaches to calculating marginal densities. *Journal of the American Statistical Association* 85:398–409.
- Geman, S. & D. Geman
1984 Stochastic relaxation, Gibbs distributions and Bayesian restoration of images. *I.E.E.E. Transactions of Pattern Analysis and Machine Intelligence* 6:721–741.

- Hurst, J. G., D. S. Neal & H. J. van Beuningen
1986 *Pottery produced and traded in North-West Europe 1350–1650*. Rotterdam Papers VI, Museum Boymans, Rotterdam.
- Jones, R. E.
1986 *Greek and Cypriot Pottery: a review of scientific studies*. British School at Athens, Fitch Laboratory Occasional Paper, Athens.
- Jones, R. E. & S. J. Vaughan
1988 A study of some "Canaanite" jar fragments from Maa-Palaeokastro by petrological and chemical analysis. in E V. Karageorghis and M. Demas (eds.) *Excavations at Maa-Palaeokastro, 1979–1986* pp. 386–398.
- Mardia, K. V., J. T. Kent & J. M. Bibby
1982 *Multivariate Analysis*, Academic Press, London.
- Pollard, A. M.
1982 A critical study of multivariate methods as applied to provenance data. *Proceedings of the 22nd Symposium on Archaeometry*:56–66.
- Pollard, A. M.
1986 Data Analysis. in R. E. Jones (ed.) *Greek and Cypriot Pottery: a review of scientific studies*, The British School at Athens, Fitch Laboratory Occasional Paper 1, Athens pp. 54–66.
- Rice, P. M. & M. E. Saffer
1982 Cluster analysis of mixed level data: pottery provenience as an example. *Journal of Archaeological Science* 9(4):395–409.
- Vitali, V.
1989 Archaeometric provenance studies: an expert system approach. *Journal of Archaeological Science* 16(4):383–391.
- Wilson, A. L.
1978 Elemental analysis of pottery in the study of its provenance: a review. *Journal of Archaeological Science* 5:219–236.

Author's address

C.E. Buck
Departments of Archaeology and Mathematics
University of Nottingham
Nottingham
GB-NG7 2RD
Email: pmzceb@vax.ccc.nott.ac.uk

