the Netherlands.
gter 1983. Seeds
nt S3 site. Final
:ohistoria 23:105-

Progress in Old World Palaeoethnobotany, Van Zeist, Wasylikowa & Behre (eds) © 1991 Balkema, Rotterdam. ISBN 90 6191 881 2

Numerical analysis in archaeobotany

4

Glynis E. M. Jones
Department of Archaeology and Prehistory, University of Sheffield, UK

ABSTRACT: The role of numerical analysis in archaeobotany, ranging from basic description to the use of multivariate statistics, is reviewed. Emphasis is placed on the archaeobotanical decisions underpinning numerical analyses. Units of observation, analysis and interpretation must be rigorously defined. Other critical archaeobotanical decisions concern the choice of data appropriate to particular questions and data sets and the level of quantification to be adopted. These decisions have major implications for the efficiency of archaeobotanical work in the laboratory. The advantages and disadvantages of pattern-searching and problem-oriented approaches to analysis are discussed. Archaeobotanists are increasingly addressing complex problems which demand multivariate analysis of large data sets. The potential of some approaches current in the related field of community ecology is surveyed. Examples are given of numerical applications in archaeobotanical interpretation and some as yet underused approaches and techniques are highlighted.

1 INTRODUCTION

Numerical analysis permeates most aspects of archaeobotany and ranges from the simple presentation of numerical data to the use of complex multivariate statistics. Clearly, archaeobotanists consider numbers fundamental to the interpretation of plant remains and it is now rare indeed to find a plant report which does not present at least the botanical identifications in some quantitative form.

This chapter is not intended as a 'recipe book' of statistical methods for archaeobotanists, but as a review of numerical analysis as an aid to the interpretation of archaeobotanical assemblages and of the implications of numerical analysis for archaeobotanical procedures. Whatever the level of expert advice sought from statisticians, archaeobotanists must be able to identify the problems and potential of their data in terms relevant to statistical analysis and to interpret the results in terms relevant to archaeobotany. Most of this chapter is, therefore, concerned with archaeobotanical, rather than statistical, decisions. These decisions can exercise considerable influence on the outcome of statistical analyses and also have a number of practical implications for archaeobotanical work in the field and laboratory.

Early attempts at quantification emphasised the assessment of economic or dietary importance. Whether based on the relative quantities of different taxa in a site/period assemblage (Helbaek 1952), on the percentage of samples in which taxa were present (Hubbard 1975, 1980) or dominant (J. Renfrew cited in C. Renfrew 1972) or on estimated dietary value (MacNeish 1967), the aim was to summarise the 'importance' of each taxon in terms of a single figure. In each case, the unjustified assumption was made that numerical values somehow reflect importance.

Instead, it is clearly necessary to make a qualitative distinction between taxa that were selected for use and those represented as discarded refuse etc. It is often assumed that species used in the recent past were also used in the more distant past, but perceived value and usage may change through time (M. Jones 1981, 1988). Dennell (1976, 1978) makes the more justified assumption that evidence for processing and storage constitutes evidence of use (lack of evidence providing no information either way). This qualitative approach, however, rests on numerical analysis, because the criteria for identifying state of processing and storage are partly quantitative.

Early attempts at quantification failed partly because they confused description (counting things) with interpretation (assessing importance). Once the unattainable goal of assigning a numerical value to the importance of taxa has been dismissed, numerical analyses are freed to address other, more answerable questions, to which end numerical description is the first step (see also Kadane 1988). The organisation of this chapter makes a clear distinction between numerical description and interpretation which is absent from much of the archaeobotanical literature on quantification.

Section 2, therefore, deals with numerical description, in particular with the type of data and level of quantification. Section 3 discusses the treatment of archaeobotanical data prior to analysis, including data selection and reduction, and problems of standardisation and transformation. Section 4 considers some of the different types of statistical analysis available (e.g. pattern-searching vs. problem-oriented, classification vs. ordination) and explores the potential for archaeobotanical interpretation of a range of techniques used in the related field of community ecology.

2 NUMERICAL DESCRIPTION

In archaeobotany, the purpose of numerical description is as a basis for the inference of past human behaviour (or occasionally for environmental reconstruction if natural seed rain can be isolated in non-charred assemblages). It is desirable, therefore, to choose for description a unit which results from a single human activity: one might call this unit the 'unit of analysis' and the activity described a 'behavioural episode'. Description at the level of the site or phase conflates the results of numerous different activities and so loses sight of the purpose of description, yet many archaeobotanical reports still offer descriptions at this level only.

A frequently chosen unit of analysis is the archaeobotanical sample and this, if restricted to a defined archaeological context, stands a relatively good chance of representing a single behavioural (or at least depositional) episode. It is often desirable, however, to sample the same context several times (in the hope of isolating separate behavioural episodes), even though these samples may ultimately prove to be of common behavioural origin. Conversely, plant remains in some contexts may be of such mixed origin that separate episodes cannot be distinguished by even the most detailed sampling (Hubbard 1976). In the former case, the same episode may be described several times, whereas in the latter the description may have little meaning because the sample is so heterogeneous, but this does not pose a serious problem until we move from sample description to interpretation of the assemblage (below, section 3.1).

Having defined the unit of analysis, consideration must be given to the choice of descriptive variables and the level of quantitative detail required. The most fundamental and widely used set of variables is the botanical composition.

2.1 Botanical composition

2.1.1 Semi-quantitative description: presence/absence

At the most basic level, taxa and plant parts may be recorded as either present or absent. This is the basis for Hubbard's presence analysis (Hubbard 1975, 1980), which is based on the concept of 'frequency' used in community ecology (e.g. Greig-Smith 1983; Kershaw & Looney 1985), and the same method is discussed by Kroll (1983) and Popper (1988) respectively under the titles of 'constancy' and 'ubiquity'. These authors summarise site/period assemblages (above, section 1), but presence/absence data can also be analysed on a sample by sample basis (e.g. Lange 1990).

A problem emphasised in the ecological literature is that presence/absence data are reliable only for samples (e.g. quadrats) of equal size since the larger the sample, the greater the chance of a taxon being present (Kershaw & Looney 1985). Not only are archaeobotanical samples often of varying size, but standardisation of samples by volume or weight of deposit does not solve the problem (below, section 3.1). Similar problems of reliability exist for more quantitative data (Orton cited in G. Jones et al. 1990), but to a lesser extent because chance presences or absences contribute relatively little to overall sample composition. Thus Lange (1990) found that the analysis of species composition was sensitive to the number of species in each sample (and so to sample volume) in the case of presence/absence data but not with fully-quantitative data.

In archaeobotany, there is the additional and common problem of pre- or post-depositional admixture from other behavioural episodes. Such contamination is particularly problematic for presence/absence data and correspondingly less so for fully quantitative data where small-scale admixtures, at least, contribute very little to total composition.

2.1.2 Semi-quantitative description: scales of abundance

A second level of semi-quantitative description involves recording of each taxon or plant part on a coarse scale of abundance (Hall et al. 1990). A major advantage of this method is that it permits rapid 'scanning' of samples (Toll 1988; Hall et al. 1990), thus saving valuable time and allowing some consideration of unproductive samples which might otherwise be left unstudied. Thus this form of semi-quantitative description provides a simple but efficient means of

(1) recording plant material (e.g. cereal bran, leaf, stem and root fragments) not easily 'counted',

(2) dealing with large numbers of behaviourally mixed samples, and

(3) selecting samples for detailed examination.

For example, urban deposits at York have produced large numbers of samples, but many are of mixed origin. Extensive scanning has provided a picture of the 'background flora' of ruderal plants growing naturally in the town and enabled routine identification of cess deposits (indicated by concentrations of cereal bran - Hall et al. 1983).

For rather different reasons, a similar method is employed in the study of storeroom deposits at Assiros. Here, large deposits of charred grain have been found, resulting from the destruction of storerooms by fire. In some instances the grain was still contained within storage jars and other facilities, but elsewhere the grain had been spilt and partly mixed with material from adjacent containers. Hundreds of samples have been collected with the aim of identifying the minimum number of separate storage entities represented. Rapid scanning of every sample is used in conjunction with stratigraphic and spatial information to 'map' the distribution of plant materials. Where mixing is relatively slight, those samples which most nearly represent the original contents of each container can be selected for detailed analysis. At the same time, more mixed areas are identified which require fully-quantitative mapping for resolution (below, section 4.3.5).

Popper (1988) uses a similar semi-quantitative method, termed 'ranking', but with rather different aims. Whereas Hall's method is essentially descriptive, no attempt being made to compensate for differential seed production or survival, Popper's method adjusts scales of abundance to take account of the expected seed production and preservation potential of each taxon or plant part. Thus Popper again confounds description and interpretation, whereas Hall leaves the problem of interpretation to a later stage. Moreover, since Popper's method of ranking is based on absolute counts rather than scanning data, it lacks the labour-saving advantage of Hall's approach.

2.1.3 Fully-quantitative description

A fully-quantitative description of each sample requires a standardised way of counting plant fragments. In other words, it is necessary to choose a 'unit of observation'. Simple counts of all plant fragments are clearly flawed (though commonplace), as these are so influenced by variable fragmentation. Moreover, while degree of fragmentation may be of interest in its own right (below, section 2.2), the recording of this information should not be confused with the basic quantitative description of taxa and plant parts.

Similar problems have been widely considered in archaeozoology, where simple counts of bone fragments have been heavily criticised because fragmentation, retrievability and identifiability are so variable. The popular alternative, however, of reducing bone fragment counts to an estimate of the minimum number of individuals represented, suffers from the serious drawback that different carcass parts are of very different utility, such that carcass parts rather than whole animals are often the most logical basic unit of analysis. In archaeobotany, estimation of minimum numbers of individuals suffers from the additional problem that the number of seeds, glumes etc. per plant is not standard in the way that the number of bones per animal usually is. Estimates of the minimum number of plants represented have been used but only in the context of reconstructing dietary importance through the use of dietary equivalents (MacNeish 1967).

At the level of description, the most useful standardised method of counting plant fragments is in terms of minimum numbers of plant parts (i.e. the smallest unit of behavioural relevance), and this is probably best achieved by a method akin to that of 'diagnostic zones' in archaeozoology (Watson 1979): for each plant part, a feature is selected for counting which (a) survives well archaeologically, (b) is unambiguously defined and (c) is accurately identifiable. To take the cereal plant as an example, embryo tips might be counted to represent grain and, in the case of glume wheats, the bases of the glumes and the culm nodes would represent chaff (rachis, glume, lemma etc.) and straw respectively; in the case of free threshing cereals, the relatively flimsy glume bases would be replaced by the more durable rachis nodes (G. Jones 1988).

2.2 Other quantitative measures

As well as botanical composition, a number of other characteristics of archaeobotanical samples can be recorded numerically. For instance, some measure of the quantity of botanical remains relative to the amount of sediment processed may be useful. This can be recorded in terms of numbers of items (e.g. seeds) per unit volume or weight of soil or, especially where large quantities of plant macrofossils are concerned, in terms of volume or weight of plant material per unit of soil. There are certain advantages to using volume of soil rather than weight. First, volume can be more easily measured on site and, secondly, it provides some means of relating the volume processed to the total volume of the deposit (Green 1979). It has been argued (e.g. Green 1979) that a standard weight or volume of soil should be processed, but this is unnecessary provided the quantity of soil processed is recorded. Indeed, since density of plant remains can vary greatly, it is often necessary to process different quantities of deposit in order to obtain the quantities of plant desirable for numerical analyses.

Density partly reflects rate of deposition and so can help to distinguish material discarded all at once from that discarded piecemeal over a period of time and mixed with other refuse (below, section 4.3). Where dung is the predominant fuel, density of charred seed may also indicate levels of fuel use and thus contribute to the identification of seasonal strata in a refuse deposit (Miller 1988).

The diversity of a sample may also be described in a number of ways, all involving the relationship between numbers of species and numbers of seeds. Thus, samples with high diversity have many species of even abundance or represented by few seeds and those with low diversity have few species of uneven abundance or represented by many seeds. The simplest expression of diversity is the ratio of number of species to number of seeds, but this index is very influenced by sample size in that larger samples do not have correspondingly larger numbers of species. Alternative measures of diversity have been applied to archaeobotanical data by Lange (1990) and Pearsall (1983) with the aim of comparing diversity between phases (and, in the former case, between feature types). Whereas Pearsall calculated a single diversity index (the Shannon index) for each phase, Lange used the mean index (Fischer's α) of the individual samples from each phase or feature type. Pearsall showed that diversity could be linked to other measures of occupation intensity, while Lange noted increasing diversity through time.

One problem with diversity indices is that samples with few species with even abundances can produce the same index as many species with uneven abundances. For this reason, Lange broke down the Shannon index into its constituent parts - number of species and evenness of abundance, showing that, while evenness increased through time, the number of species was more related to the volume of the samples. This effect could be even more marked for diversity calculated at the assemblage level (e.g. by phase), where numbers of samples as well as sample sizes could vary.

Also useful in the description of charred remains, at least, is some measure of degree of distortion and state of preservation. Hubbard and al Azm (1990) have devised numerical scales for recording the preservation and distortion of cereal grains, which can be applied at the level of the sample or of the individual grain, depending on the homogeneity of the sample. The importance of recording preservation and distortion separately for charred remains is rightly emphasised by Hubbard and al Azm on the grounds that the two result from different taphonomic processes. Distortion is the result of charring, the effects of which vary according to the severity of the charring conditions (temperature, supply of oxygen etc. - Boardman & Jones 1990) and ripeness/dryness of the grain. Preservation is largely related to post-charring conditions both before and after burial - in particular to mechanical damage and the effects of wetting and drying.

These variables may be used both as a means of identifying circumstances of charring and deposition (below, section 4.3) and as an index of the biases introduced by these processes (Boardman & Jones 1990). In addition, the roasting of milk-ripe grain (frikė) may be recognisable from the state of distortion (Hubbard & al Azm 1990).

Another characteristic which could be recorded numerically for each sample is degree of fragmentation. This could help identify such forms of processing as the pressing of olives for oil (Kyllo 1982), the pounding of wheat to make 'bulgur' (Hillman 1984a) or the coarse grinding of pulses to remove the testa and split the seed (e.g. in the manufacture of 'fava' or 'split peas' - Halstead & Jones 1989). Sample density, preservation and distortion, and degree of fragmentation by no means exhaust the range of variables which could usefully be recorded in the context of particular questions. It is impossible, however, to predict the full range of characteristics which might prove useful and it is unrealistic to record all these in every situation.

3 DATA PREPARATION

Moving from sample description to the interpretation of assemblages, some basic decisions must first be made about the form in which the data are used for statistical analysis. First, which variables should be included? This will depend partly on the questions posed, the basic approach adopted (below, section 4.1) and the type of scale on which each variable is measured, i.e. whether nominal (e.g. method of preservation - by charring, waterlogging etc.), ordinal (e.g. degree of preservation and distortion) or interval/ratio (e.g. sample density). Even for the most widely used body of data, botanical composition, certain decisions must be made and these are discussed next. It is worth noting that these decisions can have a greater effect on the results obtained than does the choice of statistical technique (Gauch 1982).

3.1 Data reduction

Ideally, each behavioural episode should be represented only once (Hubbard 1980) to avoid the risk of patterns being created by multiple sampling of the same episode or statistical tests being unreliable because extra significance is given to events represented by many samples. To minimise this risk, samples not distinguishable in terms of stratigraphy or composition should not be treated as independent units of analysis.

Once duplicate samples have been amalgamated (on paper), there are advantages to selecting for statistical analysis only the larger of the resulting composite samples, since small samples may be unrepresentative and erratic in their composition. These samples are likely to create 'noise' and obscure any pattern in the data. One strategy is to analyse small samples but to dismiss their results where they do not fit the general pattern: Lange (1990), for example, excluded samples containing few species after they had been shown to be outliers in a correspondence analysis. Since these small samples then contribute nothing to the final analysis, however, it would be advantageous to exclude them from the outset so that counting, coding and inputting of the data are avoided. If large numbers of samples are involved, considerable time may be saved in this way. The exclusion of small samples from statistical analyses does not, of course, preclude the recording of taxa not found in the larger samples.

It is more difficult to decide the threshold (in terms of quantity of plant data) at which a sample should be included in analyses. Lange (1990) compared a plot of number of species (square root transformed - section 3.3) against fraction of samples to the theoretical normal distribution and excluded samples containing too few species (three or less) to make a reasonable fit to normal. In general, these samples were also the smallest in numbers of seeds. Unfortunately, this method again involves detailed investigation of samples which are ultimately not used in analysis.

Alternatively, Fasham and Monk (1978) estimated required sample size by successive subsampling of large (pit) deposits. They took, as a measure of adequate size, the point at which fluctuations in the presence of different species, due to the addition of new subsamples, fell below 12%. Sample size was measured in terms of volume of deposit (not number of plant items) and they expressed their result in terms of the percentage rather than absolute quantity of deposit needed.

Varying rates of deposition result in very different quantities of remains per unit of deposit, however, and normally it is the quantities of taxa and plant parts (rather than volume of deposit) that are of interest. Van der Veen and Fieller (1982) have calculated the number of items required to estimate relative quantities at different levels of accuracy and confidence, for different sizes of parent population (expected numbers of items in a deposit) and for different degrees of compositional heterogeneity. With increasing size of parent population, the fraction required decreases, such that ca 400-500 items would always provide an estimate accurate to ± 5% at 95-98% confidence. Where this figure is unattainably high, samples of, say, 100 (or even 50) items could be used, provided note is taken of the resulting loss of accuracy (see for example M. Jones 1979; G. Jones

1987; Orton, cited in G. Jones et al. 1990). Whatever the level chosen, this approach has the advantage that a decision is made before too much laboratory time has been spent on material of dubious statistical worth.

As well as selecting samples for analysis, it is also advantageous to be selective of taxa and plant parts. The inclusion of rare taxa can lead to noise which may obscure patterns in the data (Gauch 1982) or even lead to misclassification of samples: Hillman (1984b), for example, notes that the presence of two seeds of a particular species resulted in the isolation of some samples as a separate group. The counter argument is that ecological 'indicator' species, though uniquely informative, are often present at low frequencies. This is undoubtedly true when dealing with live stands of vegetation, but in archaeological samples such indicator species can only be distinguished from rare occurrences attributable to minor, pre- or post-depositional contamination on the basis of regular association. The recognition of such association is only possible if the potential indicator species is present in more than one or two samples.

One solution to the problem of rarities is to group taxa into classes sharing a "common relationship with any given human activity" (Hillman 1984b), thus placing less emphasis on the isolated occurrence of a single taxon. Fewer taxa are removed in this way and less potential information is lost - particularly advantageous for 'patchy' data sets in which most species are only present in a few samples. On the other hand, the time saved by early elimination of rare taxa can be vast if a large proportion of the taxa on a site is present in very few samples and, because of their rarity, represented only by badly preserved specimens which are difficult to identify. Time thus saved can be used to study larger numbers of samples, to identify crucial specimens by time-consuming techniques such as surface-scanning electron microscopy or to collect much-needed ecological data. Early elimination of rare taxa is not, therefore, a question of 'cutting corners', but a choice between alternative archaeobotanical priorities. The onus is on the investigator to establish that the study of rare taxa makes a significant contribution to the analysis.

Whether applied to individual taxa or groups of taxa, criteria for inclusion have to be determined. Lange (1990) excluded species on the basis of the total number of seeds in the assemblage (species with less than 11 seeds were excluded). Given the potential of indicator species, however, it may be better to exclude species, not on the basis of their contribution to individual samples (since indicators are often present at low levels), but of the number of samples in which they occur. Presence in 5% or 10% of samples has often been used as the minimum in ecological studies (see for example Gauch 1982; Lange 1990). A cut-off level of 10% of samples was tested on ethnographic material from Amorgos and proved more than adequate for most questions. Indeed, for the identification of crop processing products and by-products, a cut-off level of 40% gave results which were nearly as accurate as one of 5-10% (G. Jones 1984). Different cut-off levels are appropriate for different data sets (depending on their diversity) and questions. Also, a cut-off level based on the absolute number of samples may be more appropriate than one based on the relative number.

It has also been suggested (Hillman 1984b) that only components which are conspicuously abundant under one set of conditions (e.g. in one type of crop processing by-product) and relatively rare under all others (e.g. in all other by-products) should be included in analyses. Such elimination is not necessary and may result in the loss of valuable information. For example, one class of material may be abundant in two types of crop processing by-product which can nevertheless be distinguished on the basis of a second class of material, even though this is shared with a third type of by-product. Such polythetic classification can be very useful (Clarke 1968) and is easily achieved through multivariate statistics.

3.2 Standardisation

It is most straightforward to use 'raw' data which, in the case of botanical composition, means counts of plant items. A common problem with this approach is that, if one category of samples has a much larger number of items per sample than other categories, it may be distinguished on this basis alone. For example, ethnographic samples of crop processing by-products can be distinguished on the basis of their weed seed content (below, section 4.3). Although an attempt was made to standardise the number of weed seeds in each sample (ca 300 - larger numbers would have been impracticable for most samples), the weeds in coarse sieving by-products were mostly in 'heads' and so it was necessary to collect far more than 300 seeds to ensure a representative number of heads. In a discriminant analysis using simple weed counts, the coarse sieve by-products were easily distinguished, but this was largely due to the greater sample size: artificially dividing the coarse sieving counts by three, to make them roughly comparable in size to the other samples, greatly reduced the ability of the technique to discriminate this type of by-product. Clearly the original discrimination was based not so much on the types as on the absolute quantity of weed seed, and this was simply an artefact of the method of sampling.

Green (1982) has argued for a contextual standardisation in the form of 'seed concentration analysis' which expresses each taxon as the number of seeds per unit volume of deposit. While density of plant material may illuminate rate of deposition (above, section 2.2), interpretation is far easier if sample composition and sample density are not conflated.

Assessment of sample composition is usually concerned (except for indicator species, for which presence/absence information may be useful) with relative quantities of taxa or plant parts: e.g. if light weed seeds predominate, a winnowing by-product is indicated. Compositional standardisation may, therefore, be more appropriate. Relative quantities can be expressed in terms of proportions or percentages, but with the drawback that an increase in one species always results in a decrease of all others, and one taxon present in very large numbers may obscure the relative proportions of the others. This problem is relatively slight for samples with large numbers of taxa, but worse for very small numbers. Alternatively, pairs of variables (e.g. taxa or groups of taxa) can be directly compared as ratios specifically to address certain questions. A disadvantage of proportions, percentages and ratios is that (unlike raw data) these 'composite' variables (Miller 1988) involve treating counts of different plant parts or species (perhaps with different seed productivity etc.) as equivalent. Clearly, the form in which data should be analysed depends on the nature of the data, the questions posed and the statistical methods used. Miller (1988) discusses some of the questions for which percentages or ratios have been considered particularly appropriate.

3.3 Transformation

As a rule, neither raw counts of items nor percentages are normally distributed (in the statistical sense): they are usually positively skewed (often heavily so), with many samples having few items of a particular type and few samples having many items. Many statistical procedures assume a normal or near-normal distribution and this is most crucial when tests of significance are used, though misleading results may be obtained even when description or data exploration is the main aim (Jongman et al. 1987). For some procedures that do not assume normality, it may still be desirable to give less weight to dominant species if such dominance is likely to be the result of chance events (e.g. when a single plant growing next to the sampling site produces an exceptionally high count for that species). It is often appropriate, therefore, to transform the data in a way which reduces the effect of high counts and/or makes the data more normally distributed. This is usually achieved using square roots or logarithms.

The square root transformation of ethnographic data from Amorgos and archaeobotanical data from Assiros (G. Jones 1984, 1987) reduced the skewness for all taxa, but many distributions were still far from normal. Lange (1990) used both square roots and logarithms to transform his data with apparently little difference in the result.

4 STATISTICAL ANALYSIS

4.1 Archaeobotany and community ecology

The questions posed in archaeobotany and the problems encountered have much in common with those of community ecology, a field in which statistical analysis has been more widely applied, but there are also important differences. These will be discussed in terms first of the types of interpretation sought and then of the implications for analysis.

4.1.1 Interpretation

There is considerable overlap between archaeobotany and community ecology, in that the former is frequently concerned with vegetation under human management, while the latter can provide information on the types of natural vegetation exploited. Archaeobotanical interpretation, however, emphasises past human activity, if only because samples derived from man-made contexts are most likely to shed light on this aspect of the past, whereas community ecology is more concerned with vegetation per se. Moreover, community ecology is exclusively concerned with growing vegetation (and the effects thereon of the natural and cultural environment), whereas archaeobotany also encompasses the subsequent effects of human manipulation (e.g. harvesting, food processing, discard) and of 'natural' depositional and post-depositional processes. In particular, in community ecology, spatial variation is integrally related to ecological variation in growing conditions, whereas in archaeobotany it is usually related to the final human manipulation of the plants (i.e. discard) and is only likely to reflect ecological (among other) differences when recorded at a regional level (e.g. Willerding 1980a).

Indeed, a characteristic of archaeobotanical research is concern with temporal and spatial variation, such that interpretation operates at two levels. Time trends or differences between feature types are meaningless if they cannot be interpreted in terms of past human activities. On the other hand, behavioural differences that are independent of time or space do not contribute much to our understanding of the past. Indeed, while appropriate scales can only be defined in relation to specific questions, the most interesting questions tend to be those posed on a large temporal or spatial scale (M. Jones 1985), demanding a large 'unit of interpretation'.

4.1.2 Analysis

There are two basic approaches to statistical analysis which may be termed 'pattern searching' and 'problem-oriented analysis'. In archaeobotany, pattern searching starts with counts of individual taxa, plant parts etc. (possibly standardised and/or transformed) and uses statistical techniques to group samples or to identify major axes of variation on the basis of botanical composition. These 'patterns' can then be interpreted in terms of behavioural or ecological relevance.

An advantage of this essentially inductive approach is that, by not predetermining the questions asked of the data, it allows unexpected patterns to emerge, so leading to interpretations not foreseen by the analyst. On the other hand, as applied to date, the pattern-searching approach only uses counts of individual plant taxa, implicitly excluding variables such as density of plant material, preservation and distortion, even though these might make a valuable contribution to the analysis. In any case, there will always be some potentially informative variables beyond the imagination of the analyst and, in this respect, pure pattern searching is illusory.

In community ecology, pattern-searching approaches are termed classification, where the aim is to identify discrete groups of compositionally similar samples, or indirect gradient analysis, where samples are arranged in order according to their compositional similarity (Whittaker 1973). In statistical terms, classification is achieved by various methods of cluster analysis (e.g. average link, Ward's method and the program TWINSPAN) and indirect gradient analysis by various ordination techniques (e.g. correspondence analysis or reciprocal averaging, principal components analysis and multidimensional scaling - Jongman et al. 1987). At a simpler level, correlation coefficients can be used to explore the relationship between pairs of variables. Cluster analysis and ordination operate (primarily) on compositional data and so can be easily applied in archaeobotany.

The alternative, problem-oriented analysis, begins with a specific question and uses the data to address this. In community ecology, this takes its most extreme form in direct gradient analysis where samples are arranged in sequence along a known gradient (Whittaker 1973). Regression analysis is particularly suited to direct gradient analysis; the reverse, whereby a particular gradient is inferred from compositional data, can be tackled with a variety of 'calibration' methods including inverse regression and weighted averages (Jongman et al. 1987). Analysis of variance and t-tests are appropriate when samples can be grouped into a small number of 'classes'. These are overtly deductive approaches. In an archaeological context, direct gradient analysis is only possible against a temporal or spatial (e.g. contextual) gradient, since the behavioural gradients which may underlie such trends are unknown. Calibration, on the other hand, could be used to infer behavioural gradients from compositional data.

Somewhat intermediate between direct and indirect gradient analysis is canonical ordination, which detects patterns in compositional variation that can be explained best by known gradients. Techniques include canonical correspondence analysis, redundancy analysis (the canonical version of principal components analysis) or, when samples can be grouped into a small number of classes, discriminant analysis (canonical variates analysis). Again, canonical ordination requires a known gradient (or gradients) which, in a purely archaeological context, must be either temporal or spatial. Alternatively, present-day samples (generated by known activities) can be used as a control to which archaeological samples of unknown behavioural origin are compared.

Another middle-path is to select or create variables appropriate to a particular question (preferably based on knowledge of the relevant variation in present-day samples - e.g. Hillman 1984b) and to use these 'problem-oriented' variables in essentially 'pattern-searching' statistical analyses. This may involve the construction of 'composite' variables (below, sections 4.3.2, 4.3.3) which can be a problem if heterogeneous material is involved (above, section 3.2).

The fact that interpretation is built into problem-oriented analysis can be an advantage and a disadvantage. Lange (1990) argues persuasively for the pattern-searching approach when he says that groupings of archaeobotanical remains should not be based on "notions of the investigator". He points out that the botanical contents of an archaeological feature may be unrelated to the nature or age of the feature. Similarly, behavioural (or ecological) groupings of species are inevitably based on present-day observations and the assumption that these responses have not changed through time may be unjustified in many cases. Moreover, the responses of many species are fairly catholic, embracing a range of different conditions. The grouping of samples or species in this way is seen as a "premature form of interpretation" (Lange 1990).

On the other hand, patterns in the data are predicted in advance with a problem-oriented approach and, if such patterns emerge, then a likely interpretation is that proposed in the original hypothesis or model. Such analyses are, therefore, as good (or bad) as the quality of the original model, i.e. the appropriateness of the original groupings or gradients, of modern 'control' samples and/or of the variables selected. Because their use is heuristic, however, inappropriate models are more likely to result in a failure to interpret patterns in the archaeobotanical data than in misinterpretations. The problem-oriented approach is rather less subjective than pattern searching, since it is almost always possible to interpret an observed pattern in an ad hoc fashion, but predicting a pattern to which data must be matched is a more exacting task.

Further advantages of problem-oriented analysis are noted by Hillman (1984b). First, many taxa (or plant parts) are represented in few samples and only by grouping them into classes according to their ecological or behavioural significance can reliable associations be recognised. Secondly, ratios of different classes may be indicative of certain conditions which might otherwise be masked by the mass of raw data. Thirdly, spatial and temporal variation may be dealt with in succession, thus eliminating the effects of the one in order to consider the other.

Likewise, it is desirable to allow for the effects of crop processing and discard before using samples to address more fundamental questions such as methods of crop cultivation or patterns of consumption (Dennell 1972; Hillman 1984b; G. Jones 1984, 1987). With the pattern-searching approach, on the other hand, all sources of variability are explored simultaneously, with the risk that the results of crop processing or discard may obscure (or be misinterpreted as) evidence for more interesting aspects of human behaviour.

The issues raised in the preceding discussion are illustrated in some applications of numerical analysis to archaeobotanical data (below, section 4.3) after a brief discussion of some of the statistical techniques used.

4.2 Statistical techniques

No attempt is made here to explain the mathematics underlying the statistical procedures used by various investigators. Such explanations fall outside the scope of this chapter and are, in any case, not necessary for understanding the aims and results of the analyses which follow. It may be helpful, however, to draw attention to some general points.

For interpretive purposes, a distinction can be made between 'explanatory' and 'response' variables. In community ecology, these would be the 'environmental' variables (natural and human) and the species (abundances or presences) respectively. In archaeobotany, where the data set is the archaeological assemblage, the range of explanatory variables can be expanded to include actions performed on harvested plants, depositional/post-depositional processes and time/space trends. Likewise, the range of response variables may include the different parts of plants (relevant to processing), the density of remains, their preservation and distortion (relevant to deposition/post-deposition) and so on.

It is usually desirable to treat explanatory and response variables differently in statistical analyses. In regression analysis, for example, a distinction is made between independent (explanatory) and dependent (response) variables. Canonical ordination techniques also distinguish the two types of variables. Ordination and cluster analysis deal with response variables only. Several procedures extract composite variables (gradients) which may be combinations of the explanatory variables, as in multiple regression, or theoretical gradients based

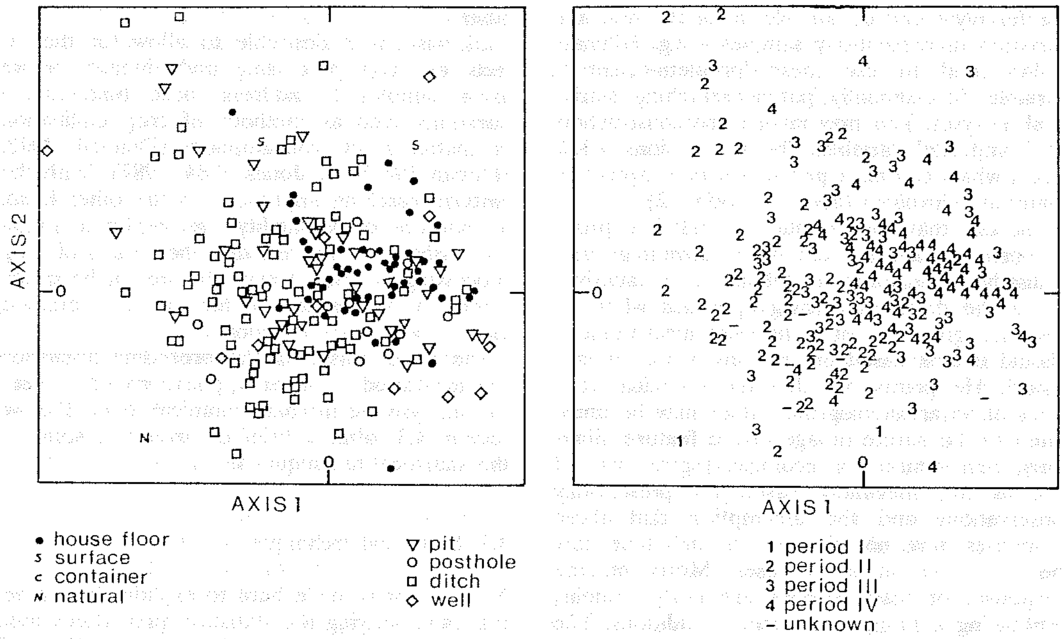

Fig. 1. Correspondence analysis of archaeobotanical samples from De Horden. Samples are represented by symbols indicating (a) main feature type (b) main period (after Lange 1990).

on the response variables, as in ordination. These composite gradients are variously referred to as ordination axes (for ordination generally), principal components, discriminant functions etc., depending on the particular technique used. In each case, the gradients serve to identify major trends in the data.

Another difference between techniques is in the number of variables that can be analysed simultaneously. For example, regression analysis handles one response variable at a time with one or (with multiple regression) several explanatory variables. Cluster analysis, ordination and canonical ordination, on the other hand, are capable of analysing many response variables simultaneously and the last also analyses one to several explanatory variables. The number of composite gradients constructed also varies, multiple regression having one and the ordination techniques several. The differences between and relative merits of different statistical techniques, as they relate to community ecology, are discussed elsewhere (e.g. Jongman et al. 1987).

A distinction should also be made between data exploration and tests of statistical significance. The former serves to summarise data or to identify patterns, while the latter state the probability that observed patterns are due to chance factors alone (i.e. sampling variability). Statistical testing, in particular, is based on the assumption that samples were selected randomly (Shennan 1988; see also M. Jones, this volume). Different statistical procedures place varying emphasis on these two aspects of statistics, with statistical testing playing a major role in univariate and bivariate analyses (e.g. analysis of variance and regression analysis) but a subsidiary role to data exploration in multivariate analyses (e.g. ordination and classification) (Gauch 1982).

4.3 Applications in archaeobotany

Methodologically, Old and New World archaeobotanists have developed parallel solutions to a number of common problems and New World archaeobotany has, therefore, featured in the preceding discussions. In line with the geographical focus of this volume, however, this section on applications of numerical analysis is restricted to Old World archaeobotany.

4.3.1 Pattern searching

Lange (1990) used the ordination technique, correspondence analysis, to explore archaeobotanical data from Late Iron Age and Roman

De Horden, the Netherlands. His analysis was based on species abundances (i.e. response or 'intrinsic' variables) and the plot of the first two ordination axes revealed a fairly homogeneous 'cloud' of samples and species points. To examine visually the effects of 'extrinsic' variables on the ordination, he replaced first the sample points and then the species points in his ordination diagram by symbols representing each of these variables in turn.

Replacing sample points first with volume of sample, then with number of species and finally with evenness of species distribution (above, section 2.2) revealed no clear ordering or clustering of these variables, indicating that the ordination is not governed by such factors. The symbols for chronological period, however, did display ordering along the first axis, indicating that this axis represents a time trend (Fig. 1). Also, the symbols for one feature type (house floors) were clustered at the 'late' end of the axis (Fig. 1), probably because house floors are almost exclusively from the last period.

Lange then replaced species points by their abundances and the resulting lack of order indicated that the ordination is not a simple function of abundance. When the species points were replaced by symbols denoting ecological group, however, cultivars and weeds were located towards one end of the second axis while grassland species and those adapted to fluctuating conditions were located towards the other end. With symbols denoting life-form, the ordering along the second axis became even clearer, presumably reflecting the same gradient from annual cereals and weeds to perennial grassland species. This axis was therefore interpreted as an 'ecological' gradient. The time trend and the ecological gradient were not perpendicular to one another in the ordination diagram: each is partly reflected in the other axis, indicating that the two trends are related.

Of particular interest is Lange's treatment of aberrant species in the ordination diagram - i.e. of species which did not relate to the ecological gradient as expected on the basis of their present-day ecology. For each such species, a separate ordination diagram was drawn showing the abundance of the species at each sample point. From this, it was apparent that some species deviated from ecological expectations because of their irregular, or very limited, distribution. Others had a split distribution, suggesting that they can favour another habitat as well as their present-day primary habitat. Yet others exhibited distributions quite different from expectations, suggesting that their ecological response has changed through time. This result could not have been obtained with a problem-oriented approach which grouped species at the outset according to their present-day ecology. A similar analysis for Late Iron Age Thorpe Thewles, N. England (van der Veen 1987), however, failed to identify any interpretable gradients.

Relationships between pairs of response variables were explored by use of correlation coefficients at Abingdon, S. England (M. Jones 1979). The proportion of cereal grain was significantly and negatively correlated with the ratio of wheat glumes to grains, implying an association between weed seeds and chaff (presumably in crop cleaning residues).

Pattern searching in the form of classification (using the program TWINSPAN) was applied to compositional data from Thorpe Thewles (van der Veen 1987) and De Horden (Lange 1990), but did not produce interpretable results, in the former case, at least, because the groups identified were based on sample size alone which affects the number of species represented (above, section 2.1.1).

4.3.2 Pattern searching using problem-oriented variables

The use of problem-oriented variables in pattern searching techniques is exemplified by Hillman's analysis of crop processing activities at Late Iron Age and Roman Cefn Graenog, N. Wales (Hillman 1984b). Principal components analysis, an ordination technique, was applied using ratios of selected plant components: the taxa and plant parts represented were classified according to (A) their probable mode of arrival on site, (B) for cereals and probable segetal weeds, the type of processing product or byproduct in which they would normally be found (on the basis of ethnographic observation) and (C) for segetal weeds, their growth habit and height in crop stands. For each sample, the numbers of items in each class were summed (calculating the crop processing 'totals' separately for cereals and weeds).

Ratios of different class totals were selected to address specific questions: e.g. the ratio of winnowing to fine sieving 'indicators' to assess whether the settlement was a primary producer of cereals or consumer of crops grown elsewhere; and the ratios of the indicators of each by-product to cleaned product indicators to identify the stages of processing represented. These ratios can then be used in principal components analyses. In the example given by

Hillman, the first principal component separates a group of 'fine cleanings' from a sample of apparently fully-cleaned grain, the variables contributing to that gradient including the ratio of waste indicators (from coarse and fine sieving) to indicators of cleaned product. The second principal component reveals two groups of samples interpreted as 'fine cleanings' and 'fine cleanings mixed with fodder or bedding', with contributory variables including the ratio of crops and segetals to species with other modes of arrival.

Variables chosen to reflect aspects of crop processing were also used in principal components analysis at Late Bronze Age Assiros, N. Greece (G. Jones 1987). For example, weed seed abundances were translated into groupings based on such characteristics as their 'aerodynamic' properties (relevant to winnowing), the likelihood of their remaining in heads or spikes after threshing (relevant to the use of a coarse sieve which allows grain to pass through it) and their size (relevant to the use of a fine sieve which retains the grain). Other variables were the ratio of weed seed to grain, the density of charred seed in the sample and the preservation and distortion of the grain.

The first principal component demonstrated an association between high density of seed, good preservation, little distortion, low ratios of weed to crop seeds, large quantities of big-free-heavy weed seeds and small quantities of small-heavy weed seeds; and vice versa. This gradient seems to distinguish between cleaned products, burnt accidentally and deposited wholesale, and by-products, burnt deliberately and discarded piecemeal as refuse. All the 'headed' weed seed groups contributed to this component, suggesting either that by-products of coarse sieving as well as fine sieving contributed to the refuse component or that some species did not remain in heads to the extent predicted. Thus, an unpredicted result was possible even though the species abundances were grouped according to predetermined criteria (cf. above, section 4.3.1).

To explore whether the variation identified by the principal components analysis was continuous or grouped, a cluster analysis (Ward's method) was performed on the same data (G. Jones 1987). Two very clear clusters were identified, the first including the samples from known storage contexts, along with some others from more ambiguous contexts. This cluster might be interpreted as cleaned products and the other, in the light of the gradient identified by the principal components analysis, as discarded by-products. When predepositional variables only (i.e. the ratio of weed to crop seeds and 'types' of weed seeds) were used, a number of sub-clusters was distinguished in the second cluster, suggesting that this discarded refuse may include a range of products and by-products.

4.3.3 Canonical ordination

The only form of canonical ordination yet applied in archaeobotany is discriminant analysis. Following an ethnographic study on Amorgos, S. Greece, four groups of present-day samples from crop processing (i.e. the by-products of winnowing, fine and coarse sieving and the cleaned product) formed the basis for the extraction of discriminant functions (G. Jones 1984). Discrimination using the proportions of individual weed species was apparently successful, in that a high proportion of samples could be reclassified into their original group on the basis of the discriminant functions extracted.

It is unlikely, however, that any archaeological site would yield exactly the same range of weed species as that found in present-day samples. To use the ethnographic samples as control groups for comparison with archaeological samples, therefore, it was necessary to translate individual species into groups based on seed characteristics relevant to crop processing. Various ways of using these characteristics were tried. Ratios of large to small, heavy to light seeds etc. were notably less successful in discriminating the four groups. To take account of the degree to which plants possess each characteristic, species were scored for each characteristic on a scale of 1 to 5. The proportion of each species was then multiplied by its score and the products summed for each sample, giving one figure for each characteristic for each sample (similar to the 'weighted averages' method of calibration in community ecology - Jongman et al. 1987). This produced a clearer discrimination, but the best result was obtained when species were grouped according to all three characteristics simultaneously to take account of the fact that one characteristic may interfere with the expression of the others. This results in six categories for each sample, ranging from big-headed-heavy seeds to small-freelight seeds. This inductive search for the best discriminating variables is justified on the grounds that, for the ethnographic samples, the 'right' answer is already known and so the problem is to find the variables which serve best to detect it.

These categories, along with the ratio of

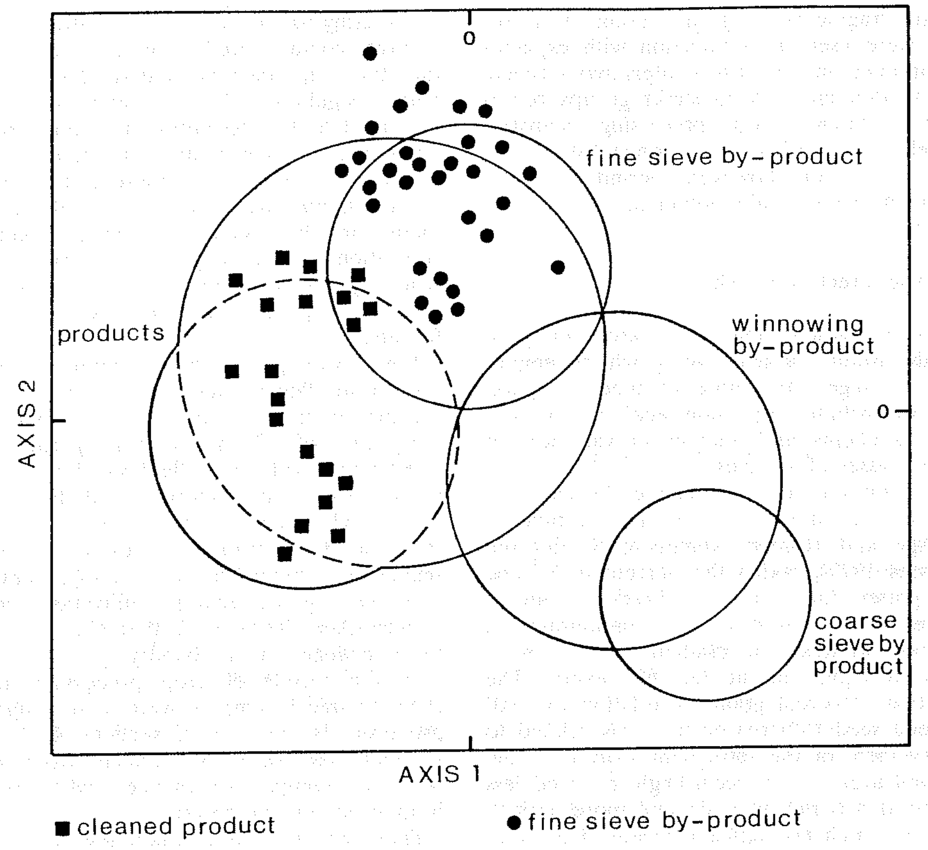

Fig. 2. Discriminant analysis of archaeobotanical samples from Assiros. Circles enclose 90% of ethnographic samples; archaeobotanical samples are represented by symbols indicating their classification (after G. Jones 1987).

weed to crop seeds, were then used as discriminating variables when archaeological samples from Assiros were compared to the ethnographic control groups (G. Jones 1987). Most of the samples were classified with high probability as fine sieve by-products or cleaned products (Fig. 2) and these classifications corresponded well with the two groups identified by the cluster analysis (above, section 4.3.2). Samples classified with lower probability could well represent mixtures of different products and by-products.

Alternatively, groups for discriminant analysis may be based on contextual criteria. For example, to test whether a contrast between 'rich' and 'poor' areas at post-medieval Amsterdam was reflected in plant usage, a discriminant analysis was performed, using species as the discriminating variables and five archaeobotanical samples from each area as the groups to be discriminated (Paap 1983). The analysis produced a clear discrimination of the two groups though the probability of samples belonging to the 'rich' groups was often low. Moreover, the discriminating species were not generally those which might be expected to reflect wealth, perhaps because of the chance effects of constructing groups from so few samples. In general, groups based on archaeological criteria are likely to be less reliable for identifying behavioural patterns than control groups of samples from known (i.e. present-day) activities, since different activities are unlikely to relate unambiguously to particular context types.

A possible non-statistical parallel for canonical ordination is Dennell's study of crop processing in Neolithic and Bronze Age Bulgaria. Response variables, in the form of compositional data (the proportion of cultigens, number of weed seeds, range of weed species, number of chaff fragments etc.) and grain measurements, were used in conjunction with explanatory variables, in this case context types (oven, floor, midden etc.), to construct groups representing different crop processing activities (Dennell 1974, 1978). This non-statistical approach does not, however, permit objective assessment of the results obtained.

4.3.4 The direct approach

For direct gradient analysis in archaeobotany, the only useful gradients along which samples can be arranged are those of time or space. Examples include regression analysis for continuous gradients and analysis of variance for discrete classes of samples.

Regression was used to analyse the changing proportions of different taxa through time at Iron Age and Roman Abingdon, S. England (M. Jones 1979), taking the percentage of one of the pottery fabrics as an indirect measure of the age of each deposit. The relationships of three taxa to this time gradient (Fig. 3) were statistically significant (at the 95% level). The proportion of cereal grain (in relation to chaff and weed seeds) decreased and was related to the movement of the settlement away from the excavated area, which accordingly received less domestic debris rich in grain and more agricultural debris rich in chaff and weeds. The other significant trends involved weed taxa: Vicia/Lathyrus species increased, suggesting depletion of soil nitrogen and so perhaps intensification of arable farming; Eleocharis decreased, suggesting either a contraction of arable away from damp ground or an improvement in drainage.

In each case, the effect of time on sample composition was interpreted in terms of changing circumstances over the period involved. Such analyses are limited to the study of one species (or other 'dependent' or response variable) at a time, though time and space could be considered together as 'independent' or explanatory variables in a multiple regression analysis.

Analysis of variance has been applied to charred remains from Anglo-Saxon Wraysbury, S. England, where the mean proportions of the major taxa were compared between different feature types: pits, gullies and ditches (G. Jones 1990). Only pulses (especially Vicia/Lathyrus) showed any significant (at the 95% level) variation between feature types, with pits and gullies having a higher proportion than ditches. T-tests showed that the proportion of pulses (especially Vicia/Lathyrus) is also significantly greater in deposits contaminated by early medieval material than in uncontaminated deposits. Since both pits/gullies and early medieval contamination tend to be concentrated in one area, this difference may be spatial, chronological, functional or a combination thereof. For example, Vicia/Lathyrus may have been used for animal fodder in the vicinity of the pits/gullies, its cultivation (presumably for fodder) may be a relatively late development or, if it was a weed of other crops, it may indicate a decline in soil fertility.

The advantage of the problem-oriented approach, in allowing questions to be tackled in a specific order, is illustrated by comparison of the ratio of Chenopodietea (garden/ruderal weeds) to Secalinetea (field weeds) in samples from present-day Amorgos and Bronze Age Assiros (G. Jones in press). T-tests showed that the ratio of Chenopodietea to Secalinetea was significantly greater at Assiros, suggesting more intensive (garden-scale) cultivation in the Bronze Age. To ensure that this pattern reflects changes in husbandry practices, rather than the results of crop processing, the archaeobotanical samples were first assigned to processing stage (above, sections 4.3.2, 4.3.3), so that fine sieve by-products from Assiros were only compared with fine sieve by-products from Amorgos and so on.

There are numerous examples of the direct approach in archaeobotany where samples are grouped according to period or feature type and then compared in terms of their numerical characteristics, but not through the use of statistical techniques.

4.3.5 Spatial mapping

In addition to spatial analysis in the sense of comparing contextual 'types', there is scope for study of the relationship between archaeobotanical material at different points in space. Spatial analysis in this sense has been applied at both the intra- and inter-site levels, but without the use of statistical mapping techniques.

At the intra-site level, within the Bronze Age storeroom complex at Assiros, individual storage entities were identified by plotting the crop composition of samples on a plan of the complex (G. Jones et al. 1986). Concentrations of certain crops were interspersed with more mixed samples indicating the centres and peripheries respectively of storage episodes. In this way an estimate was obtained of the minimum

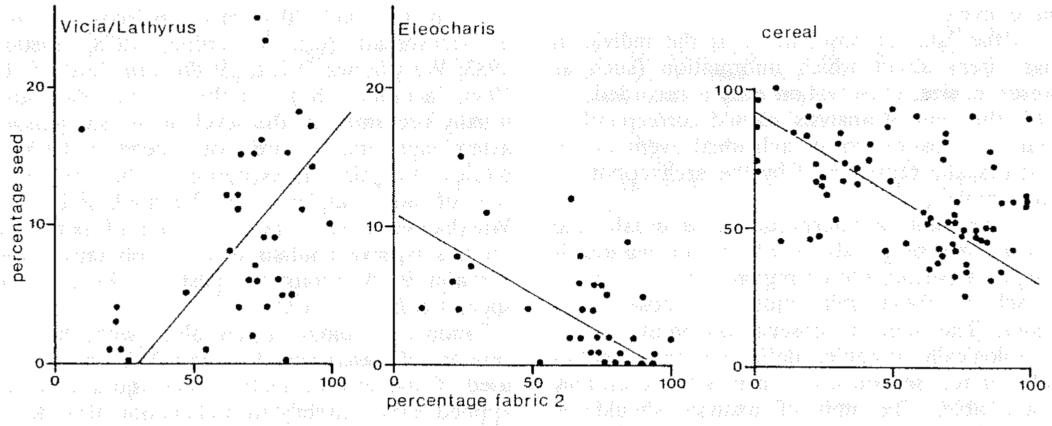

Fig. 3. Regression analysis of archaeobotanical samples from Abingdon (after M. Jones 1979).

number of storage containers of each crop in each room, which tallied well with the known number of containers.

Examples of inter-site mapping are provided by studies of Neolithic central Europe (Willerding 1980a) and, on a smaller scale, of the Iron Age Upper Thames Valley, S. England (M. Jones 1985). In the central European example, the presence or absence of a range of cultivated and wild species was plotted on a map, exposing regional differences in the adoption of crops and development of the weed flora.

In the Upper Thames Valley, the proportions of grain, chaff and weeds in each sample were compared visually for each site. Sites on the first gravel terrace have no samples with more than 50% cereal grain, while the majority of samples from the second gravel terrace had 30%-100% grain. This was interpreted as reflecting a difference between sites producing their own crops, where a considerable quantity of grain might be wasted at harvest time, and those importing crops from elsewhere, where the only waste expected would be chaff and weed seed.

5 CONCLUSIONS

The foregoing discussion has by no means exhausted the potential of numerical analysis in archaeobotany. Other areas of application include taxonomic identification, ranging from the use of metrical criteria (e.g. Helbaek 1952) to multivariate analysis (Kosina 1984) and the use of genetic data to simulate the process of cereal domestication (Hillman & Davis 1990). The focus of this chapter, however, has been on numerical analysis as an aid to the interpretation of archaeobotanical assemblages.

Large bodies of data are required to address many of the questions now posed by archaeobotanists, and in particular to distinguish variation in the archaeobotanical record which is the result of crop processing or discard activities or post-depositional factors from that caused by more interesting behavioural or ecological differences. The analysis of large assemblages poses practical problems, however, as it is difficult to compare large numbers of samples and variables on a purely visual basis. A common solution to this problem is to present and analyse data at the level of the overall assemblage, i.e. by blocks of time (periods/phases) or of space (e.g. feature types). This is unsatisfactory, however, as it obscures much variation within phases or feature types which may be of ecological, behavioural or taphonomic importance (Lange 1990). To accommodate such variation and to demonstrate the repeated associations between variables which are the basis of much interpretation, it is necessary to analyse material in smaller units such as the archaeobotanical 'sample'.

For example, many published reports use ecological or phytosociological groups in interpretation (e.g. van Zeist 1974; Wasylikowa 1978a, 1978b; Willerding 1978, 1979, 1980b, 1983; Knörzer 1984a, 1984b; Jacomet 1987) but present this information at the (site or period) assemblage level. There is a lot of scope for using this type of information in statistical analyses (above section 4.3), provided the groups are calculated at a sample level (e.g. Wasylikowa 1981; Behre 1986).

Archaeobotanical study operates, therefore, at three levels:

(1) the 'unit of observation' is the individual plant item about which information (such as presence, size, preservation etc.) is recorded;

(2) the 'unit of analysis' should correspond as nearly as possible to an individual event in the past (usually represented by the archaeobotanical 'sample');

(3) the 'unit of interpretation' is usually the larger assemblage about which information is sought (a period, site or region).

Each of these units must be chosen with rigour. The unit of observation must be archaeologically durable, definable and identifiable (above, section 2.1.3), if it is to be usefully quantifiable. The unit of analysis should be behaviourally discrete, if it is to have real meaning (above, section 2), and the unit of interpretation should be sufficiently large to provide information of archaeological relevance (above, section 4.1.1).

Consideration of the problems and potential of statistical analysis also has important implications for archaeobotanical practice. For the purposes of numerical analysis it may not be necessary or even desirable to include every sample and species or to provide fully-quantitative data for all samples. Where appropriate, data sets can be reduced in size, by excluding small samples and rare species, and in complexity, by recording species on a semi-quantitative scale. This offers the important advantage that time otherwise spent sorting, identifying and counting plant material can be used to better purpose.

A distinction has been drawn between pattern searching and problem orientation, but the approaches are complementary rather than mutually exclusive (cf. G. Jones 1987). Pattern searching may reveal unexpected patterns in the data, valuable in raising new questions or exposing unwarranted assumptions (e.g. concerning the ecology of individual species - above, section 4.3.1). Problem-oriented analysis, on the other hand, allows variability such as that caused by crop processing or discard to be controlled before exploring other aspects of predepositional behaviour (above, section 4.3.4).

It is striking how little use has been made of a number of techniques currently employed in related disciplines such as community ecology. For example, calibration is well suited to archaeobotanical questions: indicator values for different species could be used to construct gradients based on the ordering of samples, provided that such information is recorded at the sample level. In fact, reconstruction of field conditions from ecological indicator values, often in the graphical form of an 'eco-diagram', is widespread (e.g. Willerding 1978, 1980b, 1983; Wasylikowa 1978a, 1978b; van Zeist et al. 1986; Jacomet 1987), but the relevant data are usually presented at the level of the site/period assemblage and so are not amenable to numerical analysis (an exception is the presentation of two samples from Przemysl, Poland - Wasylikowa 1981). There is a lot of potential for this type of analysis in archaeobotany since the aim is to reconstruct past conditions from species composition etc.

Canonical ordination has also (with the exception of discriminant analysis) been underused. Canonical ordination techniques can be applied either directly to archaeobotanical data, using time and/or space as explanatory variables, or to modern samples collected along a known behavioural or ecological gradient, as a model with which archaeobotanical material is compared.

Archaeobotany has made remarkable advances in the sampling, retrieval and identification of plant remains, and in the development of present-day ecological and ethnographic models. Only by harnessing the potential of statistical techniques for analysing large and complex data sets, can the benefits of these advances be realised.

ACKNOWLEDGEMENTS

I am grateful to Paul Halstead, Martin Jones and Todd Whitelaw for reading and making many helpful comments on an earlier draft of this paper and to Nick Fieller and Marijke van der Veen for useful discussions on statistics generally. Any errors which remain are my own responsibility.

REFERENCES

Behre, K.-E. 1986. Analysis of botanical macro-remains. In O.van de Plassche (ed.), Sea-level research: a manual for the collection and evaluation of data, p. 413-433. Norwich.

Boardman, S. & G.Jones 1990. Experiments on the effects of charring on cereal plant components. Journal of Archaeological Science 17:1-11.

Clarke, D.L. 1968. Analytical Archaeology. London: Methuen.

Dennell, R. 1972. The interpretation of plant remains: Bulgaria. In E.Higgs (ed.), Papers in Economic Prehistory, p.149-159. Cambridge, University Press.

Dennell, R. 1974. Prehistoric crop processing activities. Journal of Archaeological Science 1:275-284.

Dennell, R.W. 1976. The economic importance of plant resources represented on archaeological sites. Journal of Archaeological Science 3:229-247.

f an 'eco-diagram',
ling 1978, 1980b,
3b; van Zeist et al.
relevant data are
of the site/period
amenable to nu-
n is the presenta-
'rzemysl, Poland -
a lot of potential
rchaeobotany since
st conditions from

ılso (with the ex-
lysis) been under-
techniques can be
ıaeobotanical data,
explanatory vari-
collected along a
ical gradient, as a
tanical material is

remarkable ad-
eval and identifica-
1 the development
and ethnographic
the potential of
ıalysing large and
benefits of these

ead, Martin Jones
ading and making
an earlier draft of
r and Marijke van
sions on statistics
emain are my own

tanical macro-remains.
Sea-level research: a
evaluation of data, p.

Experiments on the
ant components. Jour-
7:1-11.
Archaeology. London:

tion of plant remains:
ers in Economic Pre-
University Press.
p processing activities.
ce 1:275-284.
ic importance of plant
eological sites. Journal
-247.

Dennell, R.W. 1978. Early farming in south Bulgaria from the VI to the III millennia BC. British Archaeological Reports, International Series (Supplementary) 45.

Fasham, P.J. & M.A.Monk 1978. Sampling for plant remains from Iron Age pits: some results and implications. In J.F.Cherry, C.Gamble & S.Shennan (eds.), Sampling in contemporary British archaeology. British Archaeological Reports, British Series 50:263-271.

Gauch, H.G. 1982. Multivariate analysis in community ecology. Cambridge: The University Press.

Green, F.J. 1979. Collection and interpretation of botanical information from medieval urban excavations in southern England. In U.Körber-Grohne (Hrsg.), Festschrift Maria Hopf. Archaeo-Physika 8:39-55.

Green, F.J. 1982. Problems of interpreting differentially preserved plant remains from excavations of medieval urban sites. In A.Hall & H.Kenward (eds.), Environmental Archaeology in the Urban Context. London, Council for British Archaeology, Research Report 43: 40-46.

Greig-Smith, P. 1983. Quantitative plant ecology. Oxford: Blackwell Scientific Publications.

Hall, A., A.Jones & H.Kenward 1983. Cereal bran and human faecal remains from archaeological deposits. In B.Proudfoot (ed.), Site environment and economy. British Archaeological Reports, International Series (Supplementary) 173:85-103.

Hall, A.R. & H.K.Kenward 1990. Environmental evidence from the colonia: General Accident and Rangier Street. London, Council for British Archaeology, The Archaeology of York 14/6.

Halstead, P. & G.Jones 1989. Agrarian ecology in the Greek islands: time stress, scale and risk. Journal of Hellenic Studies 109:41-55.

Helbaek, H. 1952. Early crops in southern England. Proceedings of the Prehistoric Society 12:194-229.

Hillman, G. 1981. Reconstructing crop husbandry practices from charred remains of crops. In R.Mercer (ed.), Farming practice in British prehistory, p.123-162. Edinburgh, University Press.

Hillman G.C. 1984a. Traditional husbandry and processing of archaic cereals in recent times: part 1, the glume wheats. Bulletin of Sumerian Agriculture 1: 114-152.

Hillman, G.C. 1984b. Interpretation of archaeological plant remains: ethnographic models from Turkey. In W.van Zeist & W.A.Casparie (eds.), Plants and Ancient Man, p.1-41. Rotterdam, Balkema.

Hillman, G.C. & M.S.Davis 1990. Measured domestication rates in wild wheats and barley under primitive cultivation, and their archaeological implications. Journal of World Prehistory 4:157-222.

Hubbard, R.N.L.B. 1975. Assessing the botanical component of human palaeoeconomies. Bulletin of the Institute of Archaeology 12:197-205.

Hubbard, R.N.L.B. 1976. On the strength of the evidence for prehistoric crop processing activities. Journal of Archaeological Science 3:257-265.

Hubbard, R.N.L.B. 1980. Development of agriculture in Europe and the Near East: evidence from quantitative studies. Economic Botany 34:51-67.

Hubbard, R.N.L.B. & A.al Azm 1990. Quantifying preservation and distortion in carbonised seeds; and investigating the history of friké production. Journal of Archaeological Science 17:103-106.

Jacomet, S. 1987. Ackerbau, Sammelwirtschaft und Umwelt der Egolzwiler und Cortaillod-Siedlungen. Ergebnisse samenanalytischer Untersuchungen. In P.J.Suter (Hrsg.), Zürich 'Kleiner Hafner'. Berichte der Zürcher Denkmalpflege, Monographien 3:144-166.

Jones, G.E.M. 1984. Interpretation of archaeological plant remains: ethnographic models from Greece. In W.van Zeist & W.A.Casparie (eds.), Plants and Ancient Man, p.43-61. Rotterdam, Balkema.

Jones, G. 1987. A statistical approach to the archaeological identification of crop processing. Journal of Archaeological Science 14:311-323.

Jones, G. 1988. The application of present-day cereal processing studies to charred archaeobotanical remains. Circaea 6:91-96.

Jones, G. 1990. The charred plant remains. In S.Lobb & G.G.Astill, Excavations of prehistoric and Saxon deposits at Wraysbury, Berkshire. Archaeological Journal.

Jones, G. in press. Weed phytosociology and crop husbandry: identifying a contrast between ancient and modern practice. Review of Palaeobotany and Palynology.

Jones, G., V.Straker & A.Davis 1990. Early Medieval plant use and ecology in the City of London. Journal of the London and Middlesex Archaeological Society.

Jones, G., K.Wardle, P.Halstead & D.Wardle 1986. Crop storage at Assiros. Scientific American 254:96-103.

Jones, M. 1979. The plant remains. In M.Parrington, The Excavation of an Iron Age settlement, Bronze Age ring-ditches and Roman features at Ashville Trading Estate, Abingdon (Oxfordshire) 1974-76. London, Council for British Archaeology, Research Report 29: 93-110.

Jones, M. 1981. The development of crop husbandry. In M.Jones & G.Dimbleby (eds.), The environment of man: the Iron Age to the Anglo-Saxon period. British Archarological Reports, British Series 87:95-127.

Jones, M. 1985. Archaeobotany beyond subsistence reconstruction. In G.Barker & C.Gamble (eds.), Beyond domestication in prehistoric Europe: investigations in subsistence archaeology and social complexity, p.107-128. New York, Academic Press.

Jones, M. 1988. The arable field: a botanical battleground. In M.Jones (ed.) Archaeology and the flora of the British Isles, p.86-92. Oxford, University Committee for Archaeology.

Jones, M.K. 1991. Sampling in Palaeoethnobotany. In W.van Zeist, K.Wasylikowa & K.-E.Behre (eds.), Progress in Old World Palaeoethnobotany. Rotterdam, Balkema.

Jongman, R.H.G., C.J.F.ter Braak & O.F.R.van Tongeren 1987. Data analysis in community and landscape ecology. Wageningen: Pudoc.

Kadane, J.B. 1988. Possible statistical contributions to paleoethnobotany. In C.A.Hastorf & V.S.Popper (eds.), Current paleoethnobotany: analytical methods and cultural interpretations of archaeological plant remains, p.206-214. Chicago, University Press.

Kershaw, K.A. & J.H.H.Looney 1985. Quantitative and dynamic plant ecology. London: Arnold.

Knörzer, K.-H. 1984a. Pflanzensoziologische Untersuchung von subfossilen Pflanzenresten aus anthropogener Vegetation. In R.Knapp (ed.), Sampling methods and taxon analysis in vegetation science, p.250-258. The Hague, Junk.

Knörzer, K.-H. 1984b. Veränderungen der Unkrautvegetation auf rheinischen Bauerhöfen seit der Römerzeit. Bonner Jahrbücher 184:479-503.

Kosina, R. 1984. Morphology of the crease of wheat caryopses and its usability for the identification of some species - a numerical approach. In W.van Zeist

& W.A.Casparie (eds.), Plants and Ancient Man, p. 177-191. Rotterdam, Balkema.

Kroll, H. 1983. Kastanas. Ausgrabungen in einem Siedlungshügel der Bronze- und Eisenzeit Makedoniens 1975-1979. Die Pflanzenfunde. Berlin: Volker Spiess.

Kyllo, M. 1982. The botanical remains: identifications and discussion. In E.J.Peltenburg, Vrysi: a subterranean settlement in Cyprus, p.415-436. Warminster, Aris and Phillips.

Lange, A.G. 1990. Plant remains from a native settlement at the Roman frontier: De Horden near Wijk bij Duurstede. A numerical approach. Nederlandse Oudheden, Amersfoort.

MacNeish, R.S. 1967. A summary of the subsistence. In D.Byers (ed.), Environment and subsistence: prehistory of the Tehuacan Valley, vol. 1, p.290-309. Austin, University of Texas.

Miller, N.F. 1988. Ratios in paleoethnobotanical analysis. In C.A.Hastorf & V.S.Popper (eds.), Current paleoethnobotany: analytical methods and cultural interpretations of archaeological plant remains, p.72-85. Chicago, University Press.

Paap, N. 1983. Economic plants in Amsterdam: qualitative and quantitative analysis. In M.Jones (ed.), Integrating subsistence economy. British Archaeological Reports, British Series (Supplementary) 181:315-325.

Pearsall, D.M. 1983. Evaluating the stability of subsistence strategies by use of paleoethnobotanical data. Journal of Ethnobiology 3:121-137.

Popper, V.S. 1988. Selecting quantitative measurements in paleoethnobotany. In C.A.Hastorf & V.S.Popper (eds.), Current paleoethnobotany: analytical methods and cultural interpretations of archaeological plant remains, p.53-71. Chicago, University Press.

Renfrew, C. 1972. The emergence of civilisation: the Cyclades and the Aegean in the third millennium BC. London: Methuen.

Shennan, S. 1988. Quantifying archaeology. Edinburgh: The University Press.

Toll, M.S. 1988. Flotation sampling: problems and some solutions, with examples from the American Southwest. In C.A.Hastorf & V.S.Popper (eds.), Current paleoethnobotany: analytical methods and cultural interpretations of archaeological plant remains, p.36-52. Chicago, University Press.

Veen, M.van der 1987. The plant remains. In D.H.Heslop, The excavation of an Iron Age settlement at Thorpe Thewles, Cleveland, 1980-1982. London, Council for British Archaeology, Research Report 65: 93-99.

Veen, M.van der & N.Fieller 1982. Sampling seeds, Journal of Archaeological Science 92:87-98.

Wasylikowa, K. 1978a. Early and Late Medieval plant remains from Wawel Hill in Cracow (9-10th to 15th century A.D.). Berichte der Deutschen Botanischen Gesellschaft 91:107-120.

Wasylikowa, K. 1978b. Plant remains from Early and Late Medieval time found on the Wawel Hill in Cracow. Acta Palaeobotanica 19:115-200.

Wasylikowa, K. 1981. The role of fossil weeds for the study of former agriculture. Zeitschrift für Archäologie 15:11-32.

Watson, J.P.N. 1979. The estimation of the relative frequencies of mammalian species: Khirokitia 1972. Journal of Archaeological Science 6:127-37.

Whittaker, R.H. 1973. The introduction. In R.H.Whittaker (ed.), Ordination and classification of communities, p.3-6. The Hague, Junk.

Willerding, U. 1978. Paläo-ethnobotanische Befunde an mittelalterlichen Pflanzenresten aus Süd-Niedersachsen, Nord-Hessen und dem östlichen Westfalen. Berichte der Deutschen Botanischen Gesellschaft 91:129-160.

Willerding, U. 1979. Paläo-ethnobotanische Untersuchungen über die Entwicklung von Pflanzengesellschaften. In O.Wilmanns & R.Tüxen (Hrsg.), Werden und Vergehen von Pflanzengesellschaften, p.61-109. Braunschweig, Cramer.

Willerding, U. 1980a. Zum Ackerbau der Bandkeramiker. In T.Krüger & H.-G.Stephan (Hrsg.), Beiträge zur Archäologie Nordwestdeutschlands und Mitteleuropas, p.421-456. Hildesheim, August Lax.

Willerding, U. 1980b. Anbaufrüchte der Eisenzeit und des frühen Mittelalters, ihre Anbauformen, Stand--ortsverhältnisse und Erntemethoden. Abhandlungen der Akademie der Wissenschaften in Göttingen, Philologisch-Historische Klasse, 3. Folge, Nr. 116:126-196.

Willerding, U. 1983. Paläo-Ethnobotanik und Ökologie. Verhandlungen der Gesellschaft für Ökologie 11:489-503.

Zeist, W.van 1974. Palaeobotanical studies of settlement sites in the coastal area of the Netherlands. Palaeohistoria 16:223-371.

Zeist, W.van, G.J.de Roller, R.M.Palfenier-Vegter, O.H. Harsema & H.During 1986. Plant remains from Medieval sites in Drenthe, The Netherlands. Helinium 26:226-274.