Kintigh, Keith W. 1990. Intrasite Spatial Analysis: A Commentary on Major Methods. In Mathematics and Information Science in Archaeology: A Flexible Framework, edited by Albertus Voorrips. *Studies in Modern Archaeology* 3: 165-200. Holos, Bonn.

INTRASITE SPATIAL ANALYSIS: A COMMENTARY ON MAJOR METHODS

Keith W. Kintigh

INTRODUCTION

There is a large and fertile literature concerned with quantitative methods of spatial analysis in archaeology. My assessment of the state-of-the-art in spatial analysis focuses on several themes: the scale and significance of patterning, the role of models in spatial analysis, and absolute and relative indices of patterning. In this context, I hope to provide new insights on prominent techniques. While I provide a brief description of the methods and citations of the primary literature, I assume the reader to be generally familiar with them.

Before proceeding, it may be useful to step back and consider why there has not been more impact of quantitative techniques of spatial analysis on the field at large. Of the many reasons that might be offered, I focus on three: the irrelevance, the perceived impotence, and the lack of convenient access to computational procedures that can execute available methods.

Since much of the quantitative literature of intrasite spatial analysis focuses on the analysis of more or less intact living surfaces, it is *in fact* not particularly relevant to the needs of the many archaeologists whose data consist mainly of trash deposits from longer-term habitations. Important work has been done to deal with such data (notably that of Cowgill *et al.* 1984), but much more research is warranted. In this paper, however, I will not cover the analysis of these kinds of data sets.

But what of the perceived failure of the methods to deal with data from living surfaces? Other soul-searching on this topic (Clarke 1977, Kintigh and Ammerman 1982, Whallon 1984, Carr 1984, 1985, 1987) has led researchers to argue that our methods lack congruence with our substantive problems and with the archaeological record. Generally following from these arguments is the proposal of a new and purportedly better technique. Indeed, it appears that most of the novel spatial analysis techniques created in the last ten or fifteen years derived from dissatisfactions with available methods.

As mentioned above, the lack of access to computer programs to perform these analyses has been a major deterrent to the use of quantitative methods. The lack of availability of spatial analysis software has, less obviously, retarded the general level of understanding of these techniques, and as consequence, the desire to apply them. The ubiquity of microcomputers and the ease of distributing working software for these relatively standard machines has the potential to bring sophisticated methods of spatial analysis into the hands of most interested archaeologist. All analyses performed for this paper, for example, were executed on an IBM PC-compatible microcomputer with general-purpose software written by the author.

Given this emphasis during the last decade of research, where are we now? Clearly, all the basic problems are not under control. Nonetheless, I contend that not only have we made progress, but also that through thoughtful use of available techniques, it *is* possible to gain useful insights into spatial distributions of artifacts.

THE MASK SITE

The arguments presented in this paper were not developed thinking about spatial analysis in the abstract, but were formulated in the context of analyzing archaeological

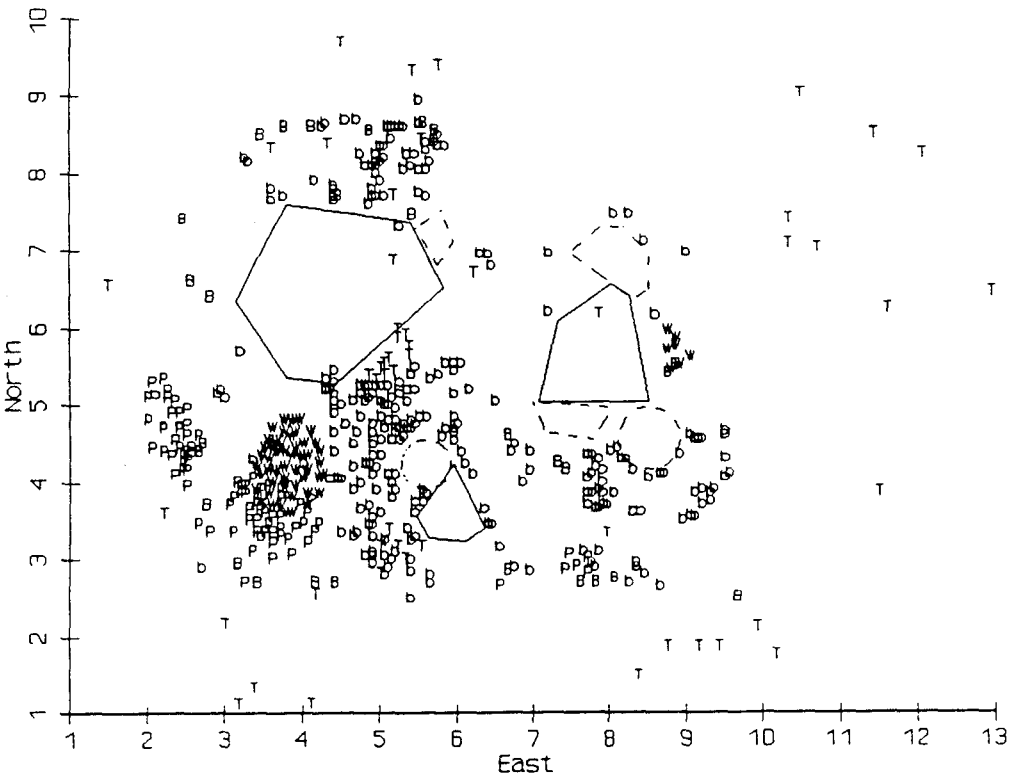

Figure 1. Distribution of artifacts at the Mask Site. T=tools, P=projectiles, W=wood scrap, B=large bone, b=bone scrap.

data. On the assumption that data-based arguments are easier to comprehend and more persuasive, I continually present illustrative results of spatial analyses. I use spatial data from the Mask Site, an Eskimo camp site (Binford 1978) that is thoroughly reported ethnoarchaeologically. In addition, Whallon (1984) has already performed and published the results of the data screening that should precede any quantitative analysis (see Whallon 1987) and has published an extensive unconstrained clustering analysis of the artifact distributions. Following Whallon, the artifact distributions will be treated as if they were from an archaeological site.

At the Mask Site, Binford recorded the exact locations of 490 artifacts spread over a 9m × 12m area. In my analyses, I use the five classes of artifacts used by Whallon: tools, projectile components (cartridge casings), wood scrap, large bone, and bone scrap. The distributions of these artifact classes are shown in Figure 1. As Whallon generously provided me with his digitized artifact locations, my analyses are based on the same data as were his.

Several points concerning the spatial patterning at the Mask Site were observed by Binford. Bone scrap was generally dropped by men seated around one of the three major hearth areas. Larger bone elements were tossed over the shoulders of these men, forming a scatter peripheral to the intensively used part of the site. Carving was localized in two areas, and target shooting in a single area away from the hearths. Tools, labeled "artifacts" by Binford, were clustered in areas peripheral to the main areas of activity, including the area near the center of the site next to boulders.

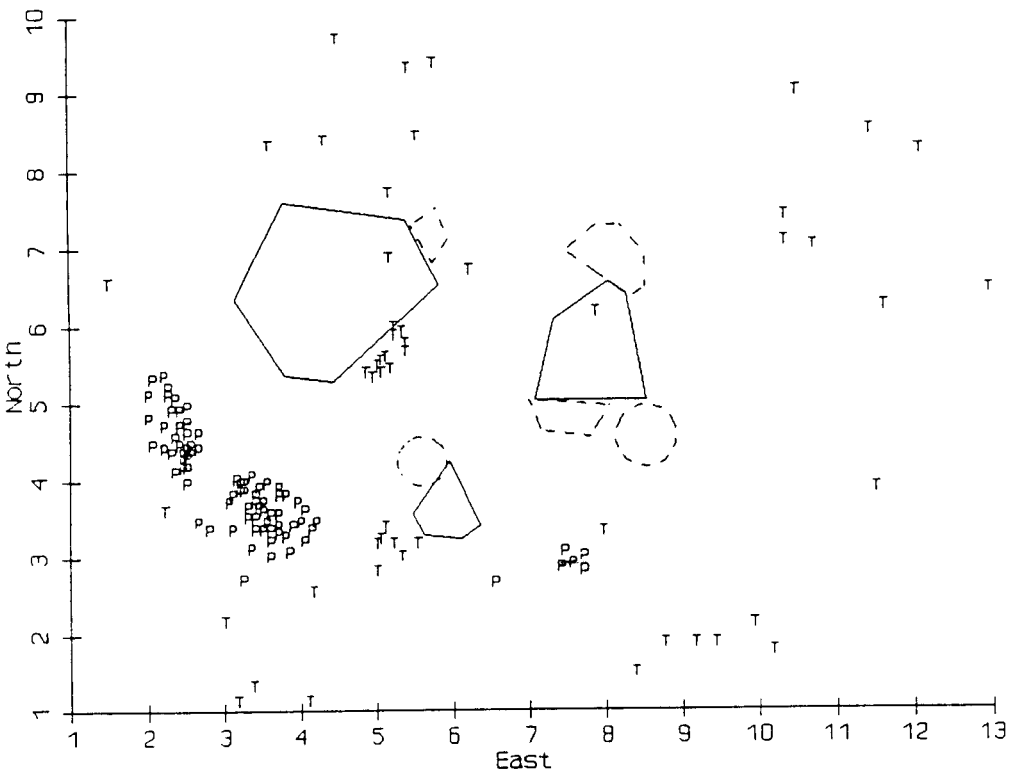

Figure 2. Mask Site tools and projectile components. R_{TP}=1.38, R_{PT}=4.85, A=0.34, C=0.39.

INDICES OF SPATIAL PATTERNING

When analyzing a spatial distribution, a question that seems often to come to mind is whether, in some absolute sense, a distribution exhibits any patterning? The logic implicit in this question is that a null hypothesis of no patterning should be examined, and that significant observable patterns may be related to the cultural or depositional processes responsible for the archaeological record. While the question seems straightforward, the answer, most assuredly, is not.

Nearest-neighbor analysis

Perhaps the most common way in which this question is operationalized is to compute the nearest-neighbor coefficient and, perhaps, test its statistical significance (Clark and Evans 1954, Whallon 1974). The nearest-neighbor coefficient *R* is defined as the ratio of the average over all points of the distances between a point and the nearest other point, divided by the average distance that would be expected if the same number of objects were distributed at random over the same area. A random distribution of points will yield a nearest-neighbor value near 1.0; a clustered distribution will yield a value less than 1.0, down to the limit of zero; and a distribution in which the objects are more evenly distributed than would be expected at random will produce a value greater than 1.0, up to the limit of about 2.15.

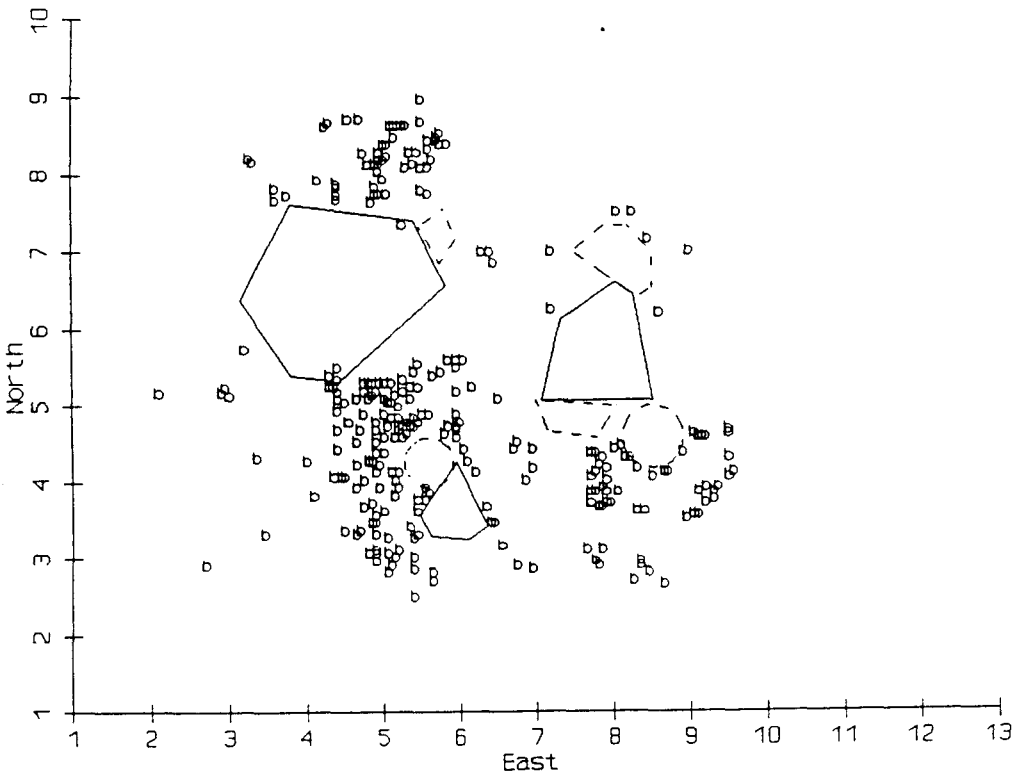

Figure 3. Mask Site bone scrap. R=0.41.

The nearest-neighbor coefficient is an *index* or indicator of the patterning of a distribution. As it is normally used, this measure is both *global* and *absolute*. It is global in that it indicates patterning in the entire distribution of points rather than the patterning in some local area within the distribution. It is absolute in that it compares the coefficient for the distribution against an absolute standard of randomness.

The boundary problem

While the nearest-neighbor coefficient appears reasonable enough, there are several serious objections to its use (Hodder and Orton 1976, Pinder *et al.* 1979). First, the nearest-neighbor coefficient is plagued with what are known as "boundary problems". One aspect of the boundary problem is that in archaeological situations it is often not altogether clear where exactly the boundary of the site is. For the Mask Site, let us assume that we have excavated the 9m × 12m area shown in Figure 1. Is the area of the site simply the 9m × 12m (108m^2) area, or might we define it more restrictively, perhaps by excluding any empty 1m × 1m grid square on the edge, an area of 65m^2?

For all 490 artifacts recorded for the Mask Site, the nearest-neighbor coefficient is 0.61 when considered within a 9m × 12m (108m^2) area, and 0.79 for the more restricted 65m^2 area. While both values indicate clustering, there is a substantial difference in the degree of clustering indicated. If the analysis is performed only on the 52 tools, together with projectiles (Figure 2), then the 108m^2 area yields a value of 0.82, indicating some clustering, whereas with the 65m^2 area the nearest-neighbor coefficient is 1.05, indicating an approximately random distribution. In many situations, especially those in which one is

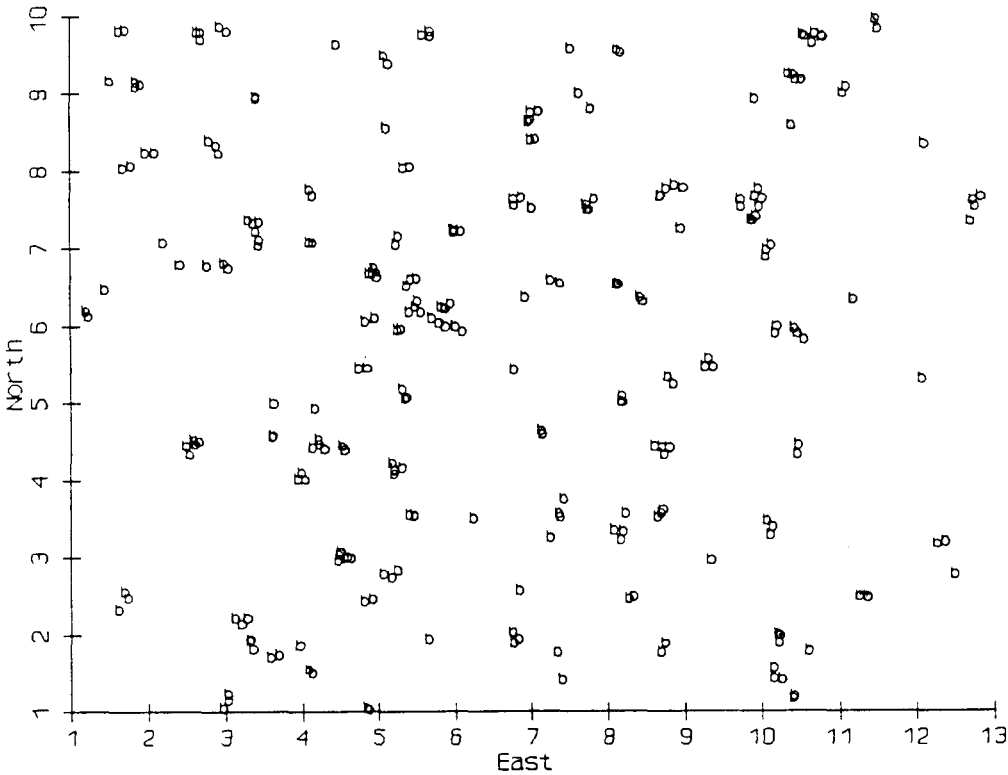

Figure 4. Artificial data with 271 points. R=0.41.

trying to deal with an entire site, the interpretation of the coefficient as an absolute measure is problematic.

The nature of spatial clusters

Another problem with nearest-neighbor analysis for assessing absolute patterning is more basic. What nearest-neighbor analysis measures as clustering is not equivalent to our intuitive notions for the spatial clustering of artifacts. The distribution of 271 pieces of bone scrap at the Mask Site shown in Figure 3 would generally be regarded as clustered, and, for the 108 m^2 area, the nearest-neighbor coefficient of 0.41 indicates strong clustering. While the distribution shown in Figure 4 has the same number of points, it does not appear strongly clustered; yet it has the same nearest-neighbor coefficient.

How is this discrepancy explained? An every-day perception of the clustering of artifacts has to do with having a relatively small number of relatively discrete groups of artifacts—the distributions of projectiles, wood (Figure 5), and bone scrap (Figure 3) are particularly clear-cut in this regard. However, the data that enter into the nearest-neighbor analysis are only the area, the number of points, and the distribution of distances from each point to its nearest neighbor. In Figure 4, the points are paired, so that each point has a relatively close neighbor, but the pairs are not frequently close to other pairs. For both the bone scrap and this artificial data, the average distance from a point to the nearest other point is 13 cm. The information that the bone scrap has three major groups of points while the artificial data has perhaps two hundred groups, is not a factor in nearest-neighbor analysis.

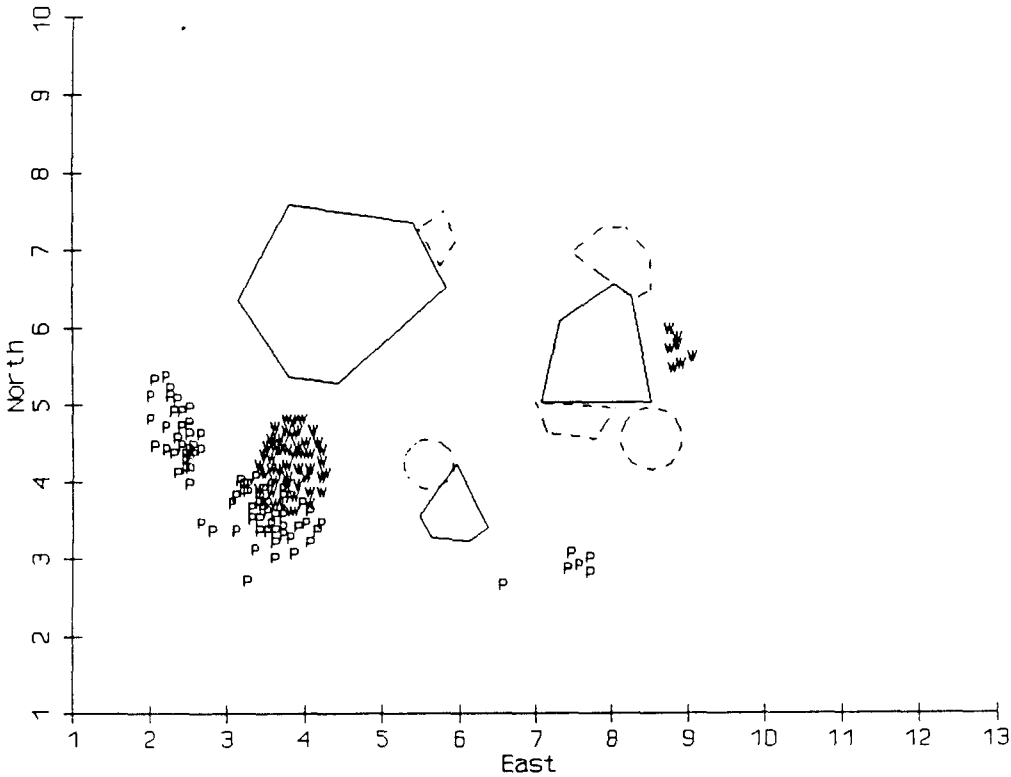

Figure 5. Mask Site projectile components and wood scrap. R_{PW}=1.15, R_{WP}=1.25, A=0.66, C=10.66.

NEAREST-NEIGHBOR COEFFICIENTS AS INDICATORS OF ABSOLUTE CLUSTERING

Although some remedies to boundary problems have been proposed (Hodder and Orton 1976, Pinder *et al.* 1979), they are not altogether satisfactory and require substantial effort to apply. The boundary problem, when coupled with the fact that the nearest-neighbor coefficient may be a substantively misleading measure of clustering, seem sufficient to say that the nearest-neighbor coefficient is *not* an appropriate absolute indicator of global patterning of artifacts.

NEAREST-NEIGHBOR COEFFICIENTS AS INDICATORS OF RELATIVE CLUSTERING

However, as others have suggested (Hivernel and Hodder 1984:100), nearest-neighbor analysis may be used cautiously as a measure of *relative* patterning of different classes of artifacts. Thus, if we examine the artifact classes separately (Table 1), the nearest-neighbor coefficients can be compared with each other. It can be shown algebraically that for boundaries with areas A_i and A_j, the nearest-neighbor coefficient for one area, $R(A_i)$, can be expressed as a simple function of the nearest-neighbor coefficient for the other area, $R(A_j)$,

$$R(A_i) = R(A_j)\sqrt{(A_j/A_i)},$$

Thus, in Table 1, the nearest-neighbor values using a 65 m^2 area are simply 1.3 times the values in the first column.

TABLE 1. NEAREST-NEIGHBOR COEFFICIENTS R FOR INDIVIDUAL ARTIFACT CLASSES.

	n	R(108 m^2)	R(65 m^2)
Tool	52	0.82	1.05
Projectile	82	0.23	0.30
Wood Scrap	57	0.16	0.21
Large Bone	28	0.43	0.56
Bone Scrap	271	0.41	0.53

While this is not particularly useful all by itself, it means that the *ratios* of the nearest-neighbor coefficients of two different tool types are the same, no matter what area is used in the calculations, as long as it includes all the points. The nearest neighbor coefficient for tools is about twice that of bone scrap, with either area. What is often the largest component of the boundary problem, the size of the area included in the site, is *not* an issue when examining the ratio of the coefficients for two artifact classes.

As a consequence, without reference to the absolute nearest-neighbor scale centered about 1.0 we can say that the distribution of projectiles, is *more* clustered than that of the tools (Figure 2), and the distribution of wood scrap is the most clustered. However, even with these relative statements, we must bear in mind that we are using the nearest-neighbor measure of clustering that is not isomorphic with the every-day meaning of clustering. While caution is required in interpreting results of this kind, nearest-neighbor analyses of separate artifact classes may well be useful *relative* indicators of the degree of global patterning displayed by different classes.

STANDARDIZING NEAREST-NEIGHBOR COMPARISONS

The foregoing discussion suggests a way to standardize the nearest-neighbor coefficients in order to compare them for different artifact classes. First, as mentioned above, one can simply look at their ratios. These ratios can be displayed in the form of a symmetric matrix (Table 2). The standardized value, then, is simply the ratio of the row variable to the column variable. For example, the neighbor coefficient of tools is 3.56 times that of projectiles, e.g., tools are much less clustered than projectiles, but the nearest-neighbor coefficients of large bone and bone scrap are approximately equal.

TABLE 2. RATIOS OF NEAREST-NEIGHBOR COEFFICIENTS, ROW NEAREST-NEIGHBOR COEFFICIENT DIVIDED BY COLUMN NEAREST-NEIGHBOR COEFFICIENT.

	N						
Tool	52	0.82	1.00	3.56	5.13	1.91	2.00
Projectile	82	0.23	0.28	1.00	1.44	0.53	0.56
Wood Scrap	57	0.16	0.20	0.70	1.00	0.37	0.39
Large Bone	28	0.43	0.52	1.87	2.69	1.00	1.05
Bone Scrap	271	0.41	0.50	1.78	2.56	0.95	1.00
			Tool	Projectile	Wood Scrap	Large Bone	Bone Scrap

The significance of nearest-neighbor results

What is the significance of these relative results? Traditional tests allow us to assess, for relatively large numbers of points, the significance of the difference between a nearest neighbor value and the random standard of 1.0 (Whallon 1974). However, I have suggested that comparisons with the absolute scale are rarely warranted. Also, one might question whether such a test is terribly meaningful, even if mathematically correct (Doran and Hodson 1975:58).

It is possible to evaluate a more useful question having to do with the significance of the relative results. Assume that the locations of artifacts at the site are taken as given. Consider the 82 projectiles (Figure 2). One might ask whether the distribution of artifacts in this class is more clustered than that of 82 locations drawn at random (without replacement) from the 490 actual artifact locations at the site. This is essentially asking whether the projectile artifact class is *significantly clustered relative to the locations of all artifacts at the site.*

In order to assess the significance of that difference using Monte Carlo methods, we need to know how frequently a value as low, or as high, as that of the actual data would be obtained by chance. Here the parameters of the distribution of *R*-values, calculated for randomized data, are of interest. The actual nearest-neighbor coefficient for projectiles, 0.23, is about 6 standard deviations below the mean of that distribution of 0.62. Thus, the nearest-neighbor clustering of the projectiles is highly significant relative to that of all artifact locations.

Spatial association of artifact classes

In previous sections, variations on nearest-neighbor analysis were used to indicate the relative degree of clustering of different artifact classes, when considered independently of other artifact class point distributions. It is also possible to use nearest-neighbor-based methods to measure the spatial association or segregation of pairs of types. Whallon (1974) proposed a graphical shared-area measure of spatial association. Hanson and Goodyear (1975) have proposed a related graphical technique based the number of shared artifacts rather than shared area (see also Clark 1979).

Use of a class-constrained variant of nearest-neighbor analysis as an indicator of class association or segregation was suggested to me by Jeffrey Parsons and is described here. In this analysis, the nearest-neighbor of a point of the first type is constrained to be a point of the second type. The expected nearest-neighbor distance is determined from the density of points in the second artifact class, which is the number of candidate points for the nearest-neighbor. The nearest-neighbor coefficient between two classes is the ratio of the mean observed nearest-neighbor distance for points of the first class, as defined above, divided by the expected nearest-neighbor distance of the points of the second class.

Table 3. Between-class nearest-neighbor coefficients. Area is 108 m^2.

	N					
Tool	52	0.82	4.85	3.51	1.46	2.82
Projectile	82	1.38	0.23	1.15	0.53	1.37
Wood Scrap	57	1.80	1.25	0.16	0.99	1.01
Large Bone	28	1.02	3.12	3.27	0.43	1.50
Bone Scrap	271	0.94	3.42	2.60	0.98	0.41
		Tool	Projectile	Wood Scrap	Large Bone	Bone Scrap

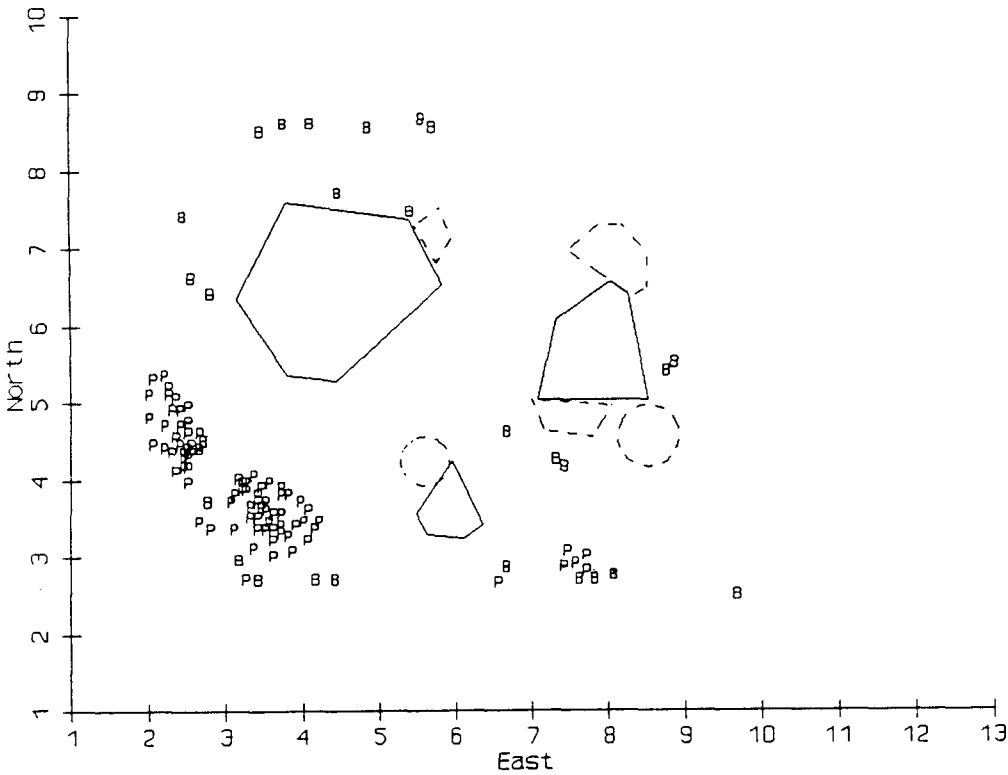

Figure 6. Mask Site projectile components and large bone. R_{PB}=0.53, R_{BP}=3.12, A=0.44, C=3.58.

Between-class nearest-neighbor coefficients are presented in Table 3. Several observations can be made about this table. First, as we would expect the diagonals of the matrix are the within-class nearest-neighbor coefficients. Second, a pair of types with values near 1.0 are interpreted as randomly intermingled, values less than 1.0 are interpreted as spatially aggregated, and values greater than 1.0 are interpreted as spatially segregated. However, the between-class nearest-neighbor coefficients do not have an upper bound as do within-class coefficients. Third, and perhaps not so obvious, the matrix is not symmetric. Consider the relationship between projectile and large bone, R_{PB} is 0.53, whereas R_{BP} is 3.12. This says that the observed distances from projectiles to large pieces of bone are about half of what would be expected given the overall number of pieces of bone, but the distance from large bones to projectiles is more than three times than would be expected given the overall number of projectiles.

These results, which may seem paradoxical, can be understood with reference to Figure 6 which shows that all projectiles have nearby bones, but many bones have no nearby projectiles. Thus, based on the first nearest-neighbors, bones might be said to be largely segregated from projectile locations, while projectiles are concentrated near large bone pieces.

Nearest-neighbor analysis, concluding remarks

As *absolute* measures of clustering, nearest-neighbor coefficients are plagued by boundary problems. However, because they can be compared with one another in ways

that are independent of the boundary of an area, nearest-neighbor coefficients can be useful measures of the *relative* patterning of different subsets of points. Using Monte Carlo methods it is possible to assess the significance of the difference between the coefficient for the actual locations of points of a class with the expected value of the coefficient under the hypothesis of a random assignment of artifact locations. Finally, the between-class nearest-neighbor coefficient can be used as an indicator of aggregation or segregation in the locations of artifacts of different classes.

However, the nearest-neighbor coefficient is still an index of *global* patterning. Because it is a global index, local clustering in one part of a spatial distribution of points may be masked by uniformity in another part of the distribution (Hietala 1984[b]:44-45). Furthermore, as the index is dependent solely on the nearest-neighbor distances, the pattern indicated may not correspond to intuitive or archaeologically meaningful notions.

Hodder and Okell's A

There seem to be fundamental inadequacies of nearest-neighbor distances for the identification of spatial clustering and spatial aggregation. The first nearest-neighbor distances of artifacts may not be good indicators of culturally interpretable clustering. Hodder and Okell (1978) and Hivernel and Hodder (1984) address this problem by formulating a measure of the spatial association of two artifact classes that uses not only the first nearest-neighbor distances, but also incorporates the distance from each point to every other point in the distribution, both of the same and the other artifact class.

Hodder and Okell's *A* is a global index that is computed directly from a matrix of intertype distances. An intertype distance is the mean distance from every point of one class to every point of the other class, presented for the Mask Site in Table 4. It should be noted that these values are simply distances, here measured in meters. Also, the matrix is symmetric, that is, the distance from tools to wood scrap is the same as the distance from wood scrap to tools. The values along the diagonal of the matrix are the distances from points of one type to points of the same type, and are the mean within-type distances. Thus, the average distance from a piece of wood scrap to another piece of wood scrap is 1.56 m, whereas the average distance from a tool to wood scrap is 4.09 m. While these numbers may have some inherent interest, they are not scaled in a way that the degree of clustering or aggregation can be inferred directly.

Hodder and Okell propose a coefficient, labeled *A*, of aggregation, or spatial association, and segregation, or spatial disassociation, of artifact classes. This coefficient for two artifact classes is defined as the product of the two mean intratype distances divided by the square of the mean intertype distance. A value of about 1.0 indicates that the distributions of the two types are spatially intermingled; the two types tend to occur and not occur in the same places. Values less than 1.0 indicate spatial segregation; the two

TABLE 4. INTERTYPE DISTANCE MATRIX FOR THE MASK SITE.

	N					
Tool	52	4.73	4.61	4.09	4.55	3.96
Projectile	82	4.61	1.54	1.91	3.80	3.60
Wood Scrap	57	4.09	1.91	1.56	3.53	3.01
Large Bone	28	4.55	3.80	3.53	4.14	3.65
Bone Scrap	271	3.96	3.60	3.01	3.65	2.94
		Tool	Projectile	Wood Scrap	Large Bone	Bone Scrap

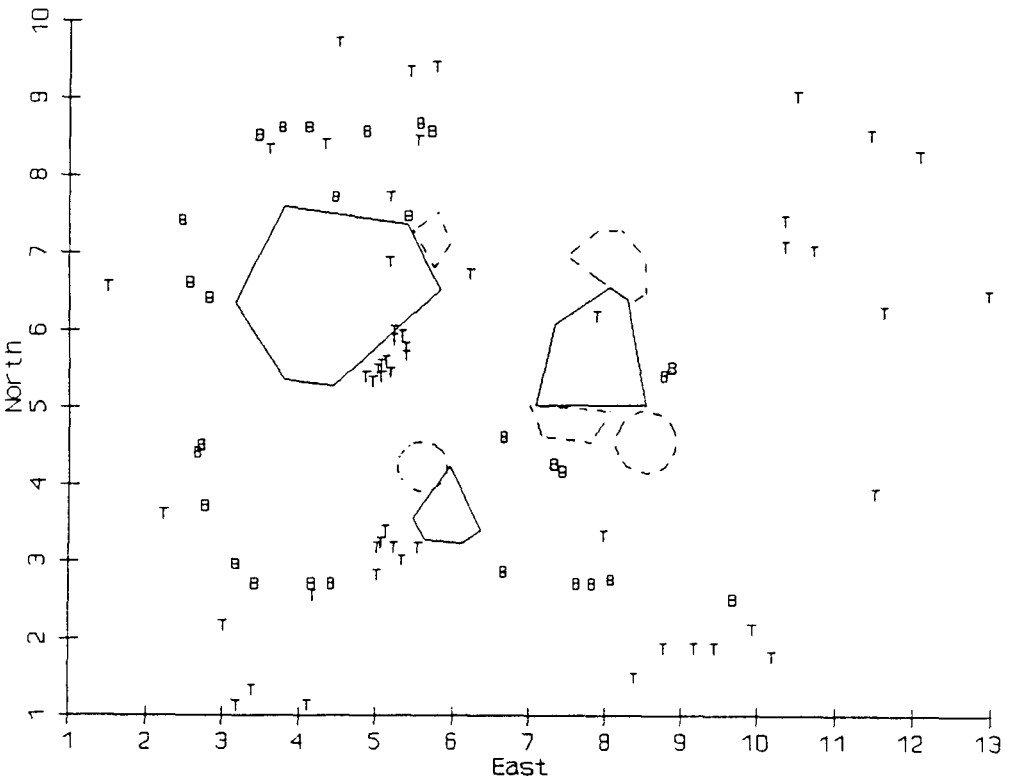

Figure 7. Mask Site tools and large bone. R_{TB}*=1.46,* R_{BT}*=1.02, A=0.95, C=1.04.*

types tend to appear in separate places; I would expect values greater than 1.0 to be empirically rare.

To understand this behavior, imagine two separate clusters each composed exclusively of a separate type. The intratype distances will be small, yielding a small numerator in the *A* ratio, whereas the intertype distances will be large, yielding a large denominator. The resulting ratio will be quite small, indicating extreme segregation. When two distributions are intermingled, the mean intratype distances for both types, that is, the average distance from a point of one type to all other points of the same type, will be about the same as the mean intertype distance producing an *A* value of about 1.0.

Comparison of Hodder and Okell's *A* and between-class nearest-neighbor coefficients

Hodder and Okell's *A* values for the Mask Site are presented in Table 5. Tools and projectiles (Figure 2) have the lowest *A*, 0.34, indicating strong segregation. Tools and large bone pieces (Figure 7) have the highest *A*, 0.95, indicating near complete intermingling. These results may be usefully compared with those obtained from the between-class nearest-neighbor analysis. The between-class nearest-neighbor coefficient between tools and projectiles is 1.38; between projectiles and tools it is 4.85. The nearest-neighbor coefficients between tools and large bone is 1.46; between large bones and tools it is 1.02. While both analyses provide intuitively reasonable results; the nearest-neighbor analysis, at least in this case, seems to me the more informative as it provides intuitively satisfying information on what amounts to two dimensions of segregation.

Significance of *A* values

Hodder and Okell have proposed significance levels for *A* for different degrees of class mixing for various sample sizes based on simulations of mixing in artificial distributions. I suggest another method. Use of the Monte Carlo method I proposed for the analysis of nearest-neighbor significance would seem appropriate for the evaluation of the significant differences from a completely intermingled situation (A=1.0). Further, this method suits the philosophy of exploiting the locational information in the distribution to reduce extraneous effects.

Evaluation of the significance of *A* requires the assessment of the likelihood of obtaining by chance values of *A* as extreme as those in an empirical situation. Using the method I propose, the artifact class assignments of the points in the original distribution are randomized with the constraint that the same number of points is assigned to each class as exist in the original distribution. While this does not change the point locations, the spatial arrangements of the *types* within the distribution are randomized, and, in general, intermingled.

The question then is, are values of *A* as low as the empirical values given in Table 5 likely to be obtained by chance arrangements of the same number of points of each type at the same locations. By repeatedly randomizing the artifact type assignments of the points, this likelihood can be assessed objectively. With 500 such randomizations, the mean value, to two decimal places, of each *A* was 1.00, and standard deviations were never greater than 0.03.

That values of *A* for the points with randomized type assignments all converge to 1.0 indicates that the observed deviations in *A* from 1.0 are not due to variations in the number of points of each type or to the patchy overall distribution of points. Further, an examination of the standard deviations from this Monte Carlo analysis shows that only one empirically observed *A*, that for tools and large bone (0.95), is within two standard deviations of the mean obtained with the randomized data. All other *A* values can be taken to differ significantly from the expected value of 1.00.

Hodder and Okell's *A*, concluding remarks

Although Hodder and Okell's *A* is a global index of clustering, it has the advantage of being an absolute measure that is *not* dependent on the boundary of an area. This measure is scale independent because it essentially takes the artifact locations as given and evaluates the degree to which the artifact types are spatially intermingled. The Monte Carlo analyses suggest that it provides a powerful measure for distinguishing segregated distributions from the intermingled ones that would be expected under conditions of random spatial location, or of close spatial association of artifact types.

Table 5. Hodder and Okell's A for the Mask Site.

	N					
Tool	52	1.00	0.34	0.44	0.95	0.89
Projectile	82	0.34	1.00	0.66	0.44	0.35
Wood Scrap	57	0.44	0.66	1.00	0.52	0.51
Large Bone	28	0.95	0.44	0.52	1.00	0.91
Bone Scrap	271	0.89	0.35	0.51	0.91	1.00
		Tool	Projectile	Wood Scrap	Large Bone	Bone Scrap

Hodder and Okell's *A* and the between-class nearest-neighbor coefficient occupy opposite ends of a continuum. *A* is a function of all neighbor distances, while the between-class nearest neighbor coefficient is based only on first nearest-neighbor distances. Both may provide useful information about a distribution, but neither seems to be a sufficient measure of spatial association.

Local density analysis

Local density analysis, proposed by Johnson (1976, 1984; also discussed by Graham 1980) provides an absolute, global measure of artifact class association that considers interpoint distances only within a fixed radius of each point, indicating patterning *at a certain scale*. The local density coefficient $C_{ij(r)}$ is the mean density of points of type j in the neighborhood of points of type i, divided by the global density of type j points. To calculate this measure, each point of type i is considered in turn. The local density of points of type j for a given i point is the number of points, M_{ij}, of type j that are found within a fixed radius r of the subject point divided by the area of the neighborhood, πr^2. The global density of type j points is simply the total number of type j points (N_j) divided by the area of the site (A).

For the Mask Site the symmetric matrix of local density coefficients for a 1m radius is presented in Table 6. Local density coefficient values about 1.0, such as between tools and large bone (Figure 7, C=1.04) indicate that there is no aggregation or segregation of the types at this scale of analysis, i.e., the average actual and expected densities are about the same. For the Mask Site, the average number of large bones found within a meter of each tool is about the same as the number one would expect if the large bones were randomly distributed.

A value greater than 1.0 indicates spatial association of the two types. For projectiles and large bones (Figure 6), the local density coefficient of 3.58 indicates that an average of more than three times as many large bones are found in the neighborhood of projectiles as would be expected at random. Local density coefficients less than 1.0 indicate segregation of the types. Mask Site tools and projectiles (Figure 2), for which C=0.39, are highly segregated. Indeed, only about 40% of the number of projectiles are found in the 1.0m radius neighborhood of tools. The local density coefficient of a point type with itself is similarly interpreted. Except for very small numbers of points or situations in which boundary problems, to be discussed below, come into play, values of intratype coefficients substantially less than 1.0 should not occur.

BOUNDARY PROBLEMS WITH THE LOCAL DENSITY COEFFICIENT

Because the local density coefficient is a global measure that is dependent on the site area, it can be expected to have some boundary problems. To the extent that fixed-radius

TABLE 6. LOCAL DENSITY COEFFICIENTS A NEIGHBORHOOD RADIUS OF 1.0 M FOR THE MASK SITE.

	N					
Tool	52	2.87	0.39	0.01	1.04	2.63
Projectile	82	0.39	15.22	10.78	3.58	0.86
Wood Scrap	57	0.01	10.78	25.46	1.27	2.32
Large Bone	28	1.04	3.58	1.27	2.46	1.78
Bone Scrap	271	2.63	0.86	2.32	1.78	5.28
		Tool	Projectile	Wood Scrap	Large Bone	Bone Scrap

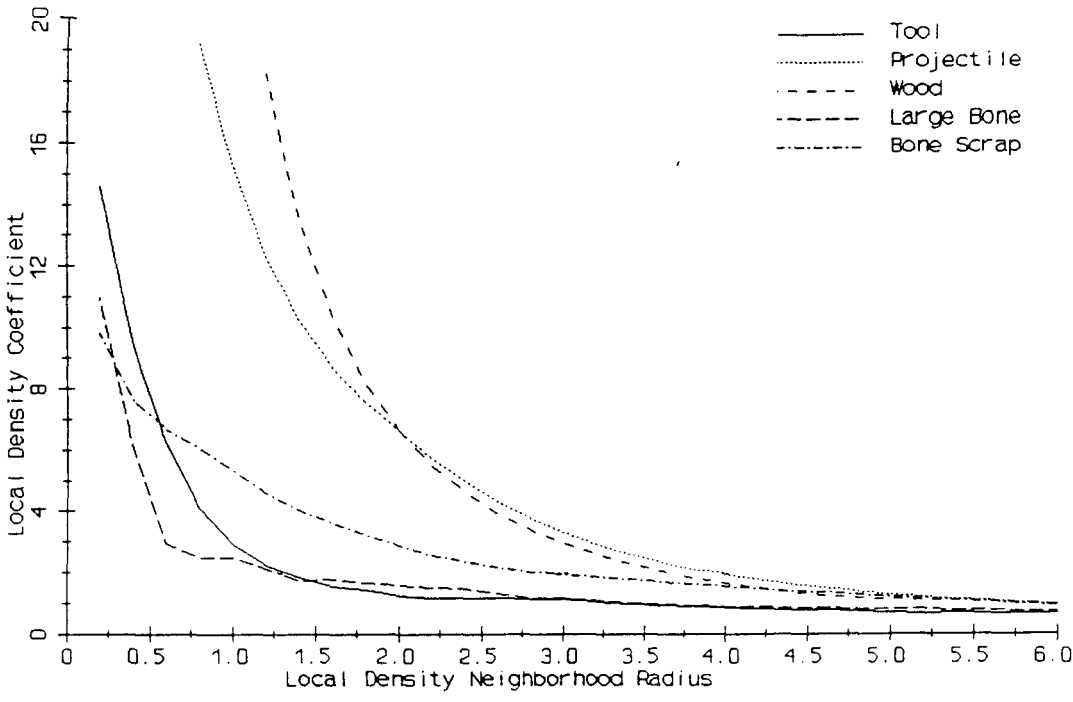

Figure 8. Mask Site intratype local density plot.

neighborhoods around points extend beyond the boundary of the site, the local density coefficient will be numerically depressed, because a global density computed on the site area will be too high. If the site boundary is enlarged to accommodate all points plus their neighborhoods, the global density will be unrealistically low and the local density coefficients will be elevated.

Local density analysis as a relative measure

A useful approach to the boundary problem is to view the local density coefficient in a *relative* framework, much as was suggested above for nearest-neighbor analysis. It is easily shown that, for a given neighborhood radius r, the local density coefficient computed for one area is related in a simple way to the coefficient computed using a second area and it can be shown that the ratio of two local density coefficients is independent of the site area. Thus, one can compare coefficients to one another without regard for the site area. For a 1m radius neighborhood, projectiles and large bones (Figure 6, C=3.58) are more strongly aggregated than are tools and large bone (Figure 7, C=1.04).

Significance of local density values

The Monte Carlo method described above can be used to derive significance levels for local density coefficients. That is, taking the point locations as given, but repeatedly randomizing the assignment of point types to locations will provide an assessment of the likelihood of obtaining, by chance, a coefficient at least as extreme as the actual value. While this analysis has not been fully carried out, limited Monte Carlo experimentation

hints at an interesting result. It appears that randomized point type assignments result in *all* local density coefficients in the matrix tending to the local density coefficient obtained in an analysis in which the entire distribution of points is considered as being of one type. This result seems sensible and is directly analogous to the results of the Monte Carlo analysis of the within-class nearest-neighbors.

Local density analysis and patterning at different scales

Because the local density coefficient is dependent on the specification of the neighborhood radius, it allows us to examine patterning at a specific scale. In a sense, it occupies an intermediate position between a nearest-neighbor analysis and Hodder and Okell's *A* which amounts to an all-neighbor analysis. This seems advantageous, as *behavioral* notions of spatial patterning are scale dependent.

Examining the results of local density analyses performed for a range of neighborhood radii can reveal information about the scale, or scales, of patterning of artifact classes or pairs of classes. Because each analysis yields an entire matrix of coefficients, a convenient method of evaluating the information is needed. One obvious method is to plot the coefficients against the neighborhood radii. For example, Figure 8 shows the intratype coefficients for the Mask Site plotted against neighborhood radii from 0.2m to 6m. In these plots, local density coefficients greater than the maximum values on the vertical axes are not shown. The objective of examining these plots is to compare behavior of the coefficients for the different artifact class combinations and to identify critical radii.

Graham (1980:110) illustrates expected shapes for these plots in three ideal cases. However, analysis of the Mask Site data and artificial data indicate that the interpretation of these plots is less straightforward than Graham's discussion suggests. A distribution in which most points are concentrated in a single dense cluster, such as Mask Site projectiles or wood (Figure 5), shows high local density values at the small radii that fall off to a value near 1.0 as the radius increases (Figure 8). Because these types are so clustered, even at very small radii the average number of artifacts of the same type in the neighborhood is extremely high relative to the expected number.

For wood, increasing radius from about 1.5m up to about 5m includes more area, and hence a greater expected number of pieces without increasing the observed number of artifacts within the neighborhood at all, so the local density coefficient gradually decreases. For projectiles, however, as the radius increases from about 1m to 3m, the neighborhoods around the projectiles in each of the two large clusters incorporate projectiles from the other cluster, so the coefficient declines less rapidly than that of the wood.

Further examination of Figure 8 shows that *all* the intratype coefficients have the same general shape. They reach their highest values at the smallest radius computed and then decrease as the radius increases. The local density coefficient for bone scrap (Figure 3), which is concentrated in large and rather more sparse clusters, is only moderately high for small neighborhoods, and decreases only slowly over the range of radii. While shape alone is not a clear-cut indicator of the kind of clustering exhibited by a set of points, it is notable that projectiles and wood, clearly the most clustered of the types, have the highest local density coefficient across the range of radii considered.

Graham indicates that several widely separated and evenly spaced clusters result in curves with local density coefficients that are high at small radii which correspond to the small clusters, then dip below 1.0 at intermediate radii, then increase again as the neighborhood radii begin to include points in other clusters, and then gradually decline again to near 1.0 at large radii. Figure 9 and Figure 10 show two artificial distributions of points that were generated in an identical manner, but with different densities. Both are strongly clustered, yet their intratype local density coefficient plots have different shapes (Figure 11), and the shapes are markedly different from the ideal suggested by Graham.

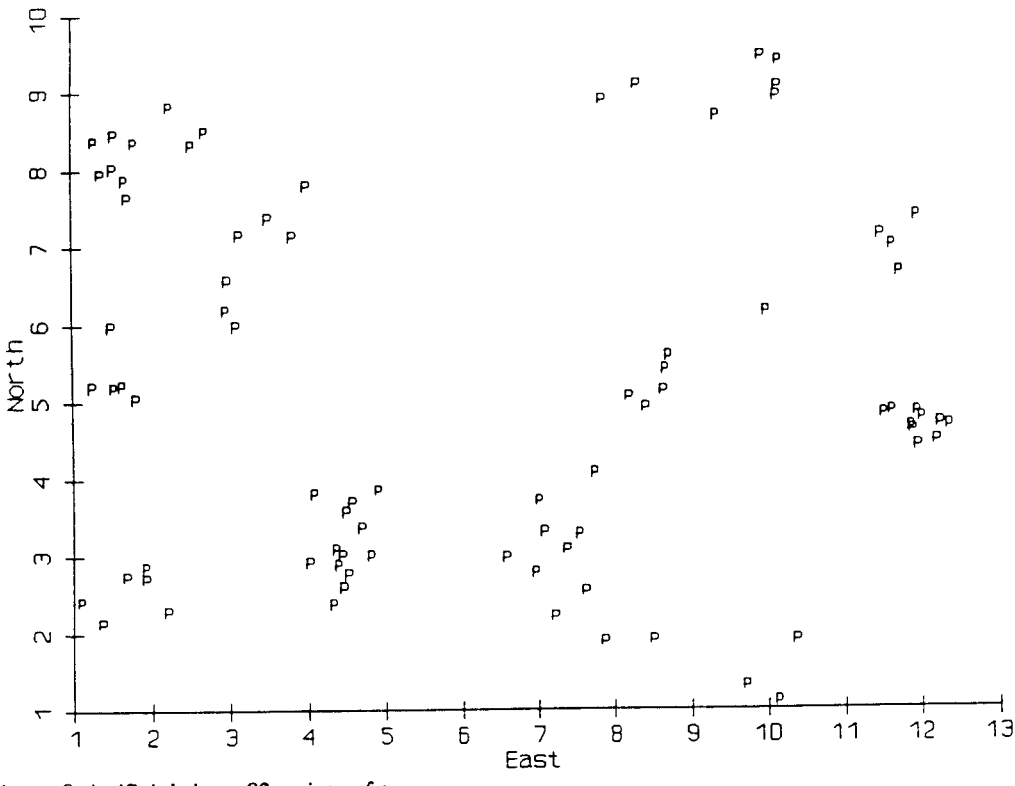

Figure 9. Artificial data, 82 points of type P.

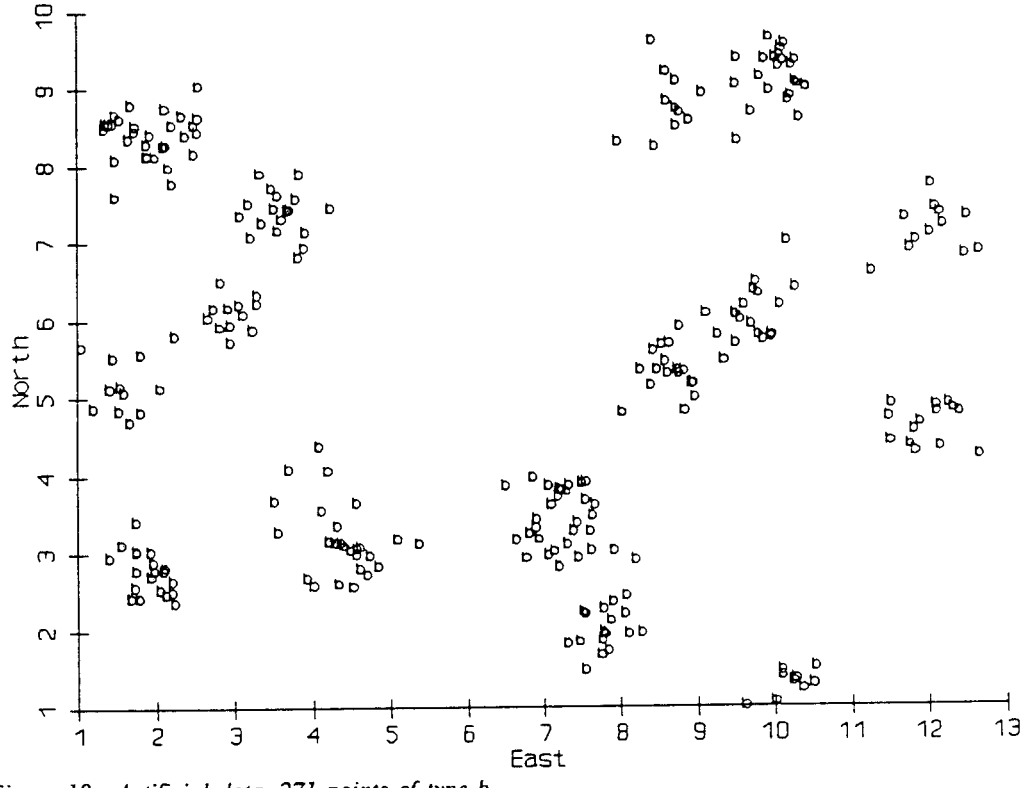

Figure 10. Artificial data, 271 points of type b.

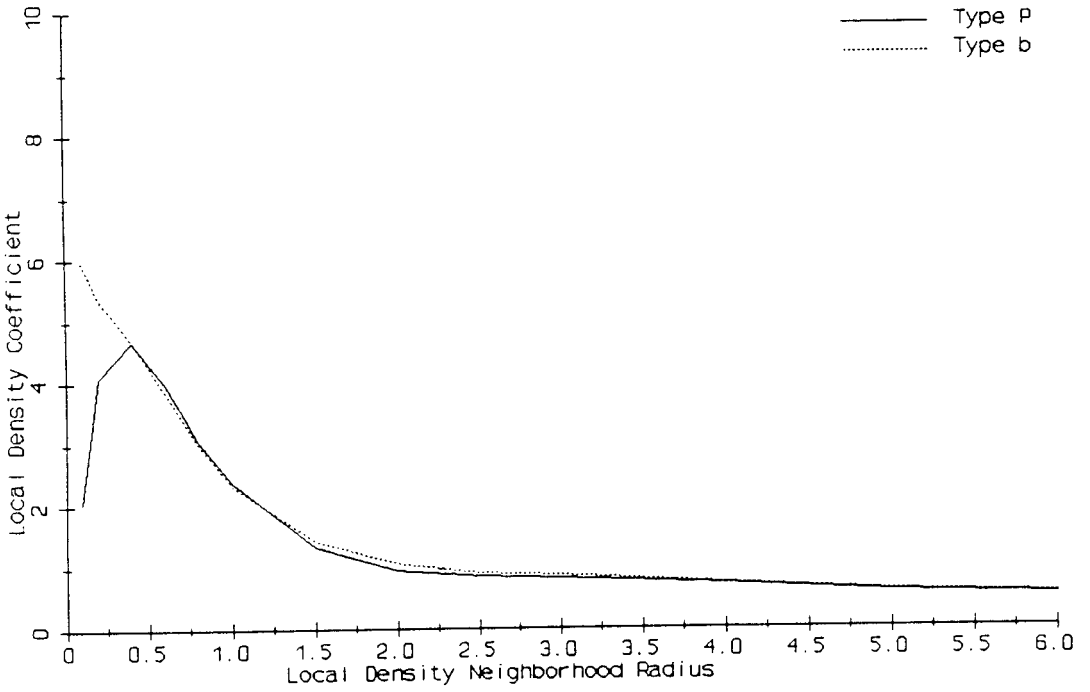

Figure 11. Local density plot for artificial distributions of type P and b points.

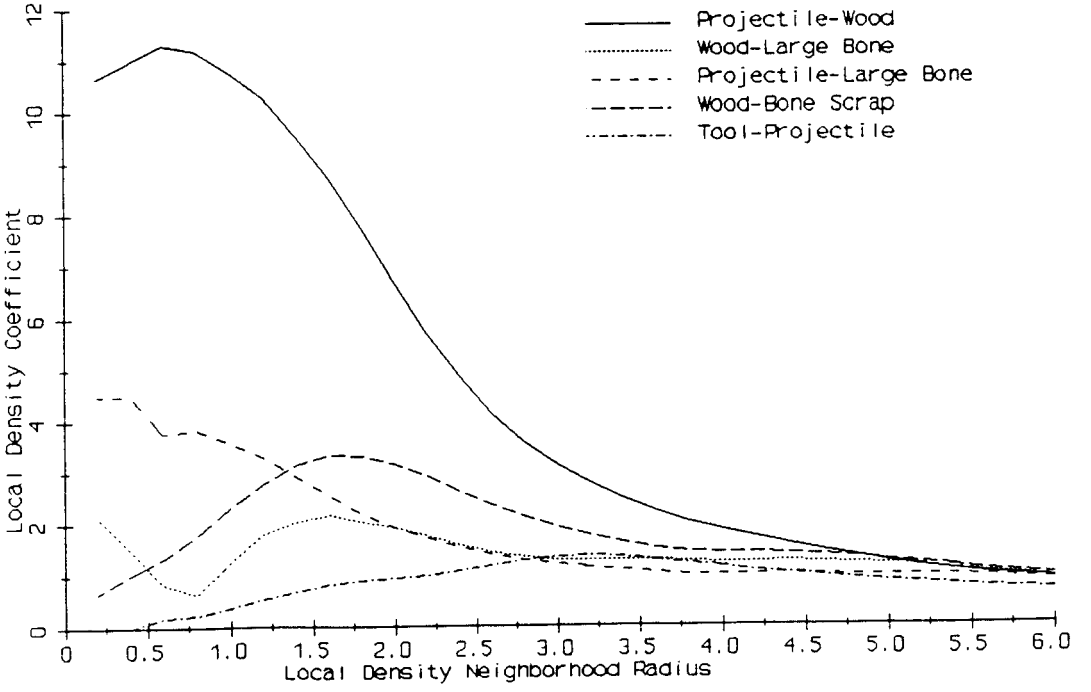

Figure 12. Mask Site intertype local density plot.

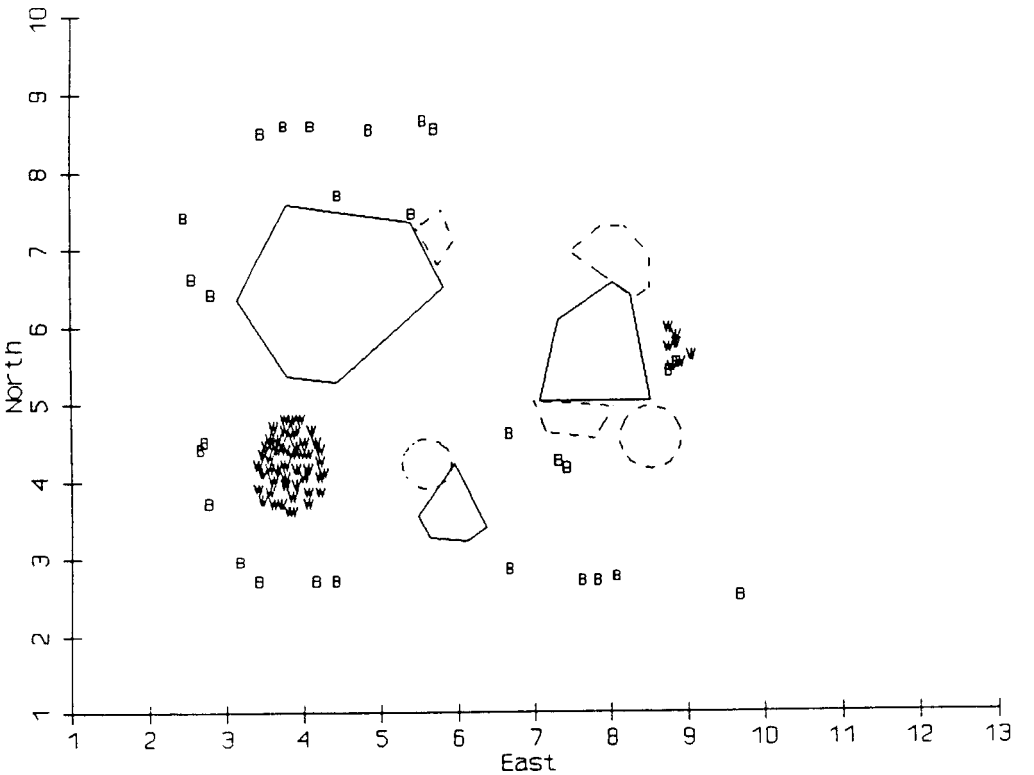

Figure 13. Mask Site wood scrap and large bone. R_{wB}*=0.99,* R_{Bw}*=3.27, A=0.52, C=1.27.*

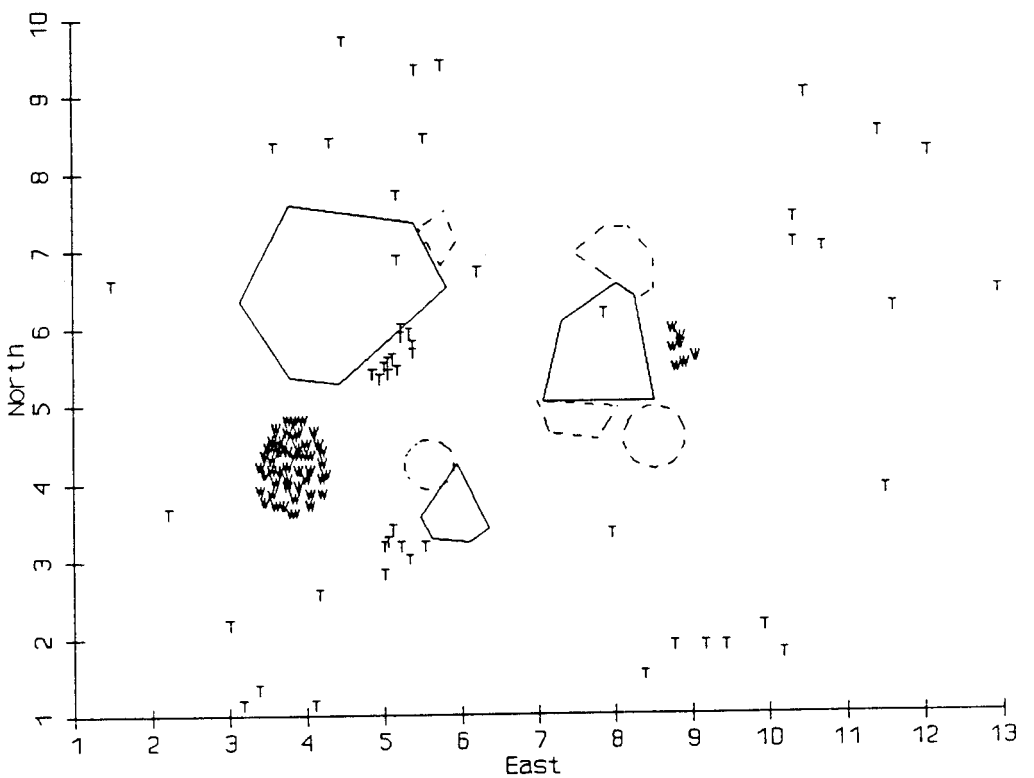

Figure 14. Mask Site tools and wood scrap. R_{TW}*=3.51,* R_{WT}*=1.80, A=0.44, C=0.01.*

The intertype coefficients display a different set of behaviors. The most obvious aspect of Figure 12, which displays all intertype coefficients for the Mask Site, is that over almost the entire range of radii, the projectile-wood local density coefficient is very much greater than any other. Examination of Figure 5 shows that both types appear in dense clusters, and that a large fraction of points of each type is within 1m or 2m of a large portion of points of the other type.

Another interesting pattern is shown by the wood-large bone (Figures 12, 13) coefficient, which is relatively high, about 2.0, at a radius of 0.2 m, drops quite low, below 1.0, at a radius of 0.8 m, rises above 2.0 again at a radius of about 1.6 m, and then gradually falls off. This unusual behavior seems to be due to a tight concentration of several pieces of wood and large bones and by a ring of large bone pieces about a meter distant from a dense cluster of wood pieces.

Comparison of local density analysis, Hodder and Okell's *A*, and nearest-neighbor analysis

Like Hodder and Okell's *A*, but unlike the between-class nearest-neighbor coefficient, the local density coefficient between artifact types is symmetric. Thus, only a summary measure of aggregation or segregation can be obtained. It is not possible to distinguish asymmetric associations from ones in which the spatial associations are mutual. For example, projectiles all have nearby bones, but many bones do not have nearby projectiles (Figure 6). The ambiguity of the situation is highlighted by the contradictory indications provided by local density analysis and Hodder and Okell's *A*. In this example, the local density C_{PB}=3.58 indicating aggregation whereas Hodder and Okell's A_{PB}=0.44, suggesting segregation of these types. Between-class nearest-neighbor analysis elucidates the actual situation: R_{PB}=0.53, whereas R_{BP}=3.12; projectiles are universally close to large bones, but not vice versa. Although Graham (1980:108) considers the asymmetry of between-class nearest-neighbor analysis to be a liability for his purposes, it seem here to be an important asset.

A local density analysis (r=1.0) and between-class nearest-neighbor analysis of obviously spatially segregated tools and wood (Figure 14) indicate spatial segregation (C=0.01, R_{TW}=3.51, and R_{WT}=1.80). While this local density analysis indicates extreme segregation, Hodder and Okell's *A* of 0.44 indicates spatial aggregation of the same magnitude as is found between projectiles and large bone (Figure 6).

Conclusions: indices of spatial patterning

Nearest-neighbor analysis, Hodder and Okell's *A*, and local density analysis, with the extensions suggested here, may all provide useful and somewhat different information concerning spatial patterning. Two major questions will be considered in turn.

Within-class clustering.

Nearest-neighbor analysis and local density analysis can provide *global* indications of the *relative* clustering of artifact classes. The only intra-class information provided by the computations performed for Hodder and Okell's *A* is the mean within-class interpoint distance, which is not directly indicative of clustering. A low value of the nearest-neighbor coefficient does indicate clustering, it is, however, not a reliable indicator of behaviorally meaningful clustering because the analysis considers only the first nearest neighbors. The maxima shown in local density coefficient falloff curves provide an index to the degree and a rough idea of the spatial scale of clustering, and the rate of falloff indirectly

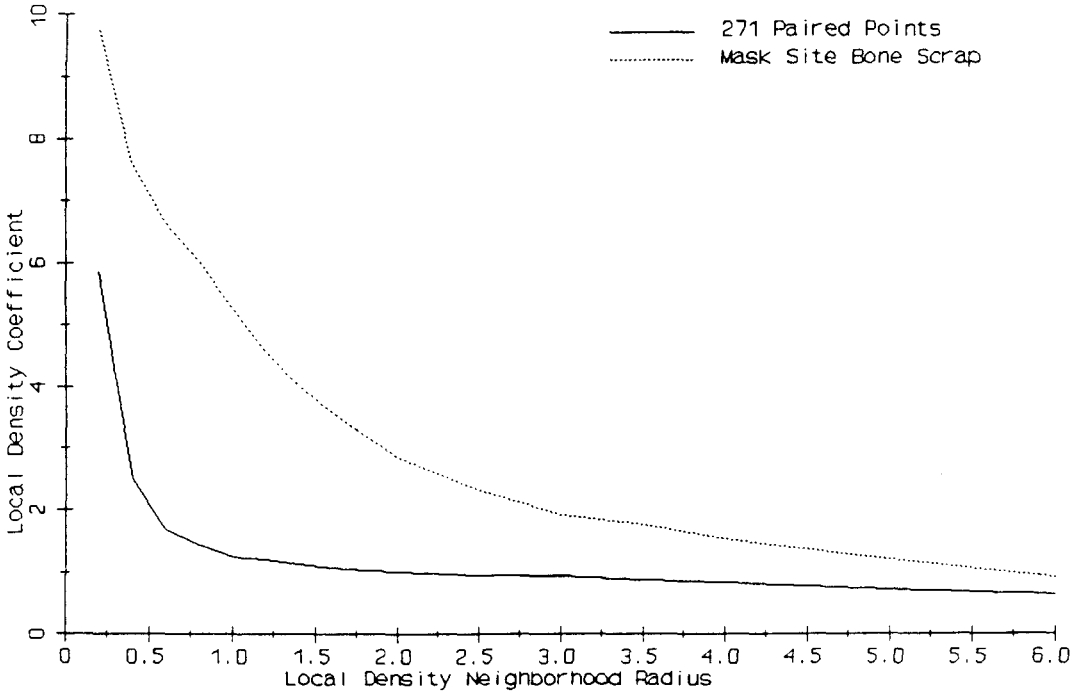

Figure 15. Local density plot for artificial distribution of 271 points shown in Figure 4 and Mask Site bone scrap (Figure 3).

indicates the spread of the clusters, but, usually, no clear evidence for the number of clusters. However, local density analysis does clearly distinguish the cases with identical nearest-neighbor results shown in Figures 3 and 4 (Figure 15).

Spatial aggregation and segregation.

All three methods attempt to deal with this question. Although Hodder and Okell's *A* provides an absolute measure, its inclusion of all interpoint distances seems apt to cloud or obscure much behaviorally relevant patterning.

As a scale-sensitive measure, local density analysis appears to have an inherent advantage over nearest-neighbor analysis. Local density falloff plots may indicate relative global patterning of different combinations of artifact classes at different spatial scales. While interpretation of the plots is not direct, aggregation or segregation of artifact classes at particular spatial scales may be highlighted. This may be a major advantage, especially when large numbers of artifact classes must be dealt with.

Despite its liabilities, between-class nearest-neighbor analysis, used as a relative global measure of aggregation or segregation, can indicate asymmetric relationships that I suspect are common in archaeological data.

PURE LOCATIONAL CLUSTERING

Pure locational clustering refers to a method first proposed by Kintigh and Ammerman (1982) that, due to our unfortunate failure to name it, came to be known as *k*-means spatial analysis. Pure locational clustering is a rigorous, numerical procedure that was

intended to approximate the results of an intuitive division of a point distribution into a set of clusters. That is, its principal objective is *not* to provide an index for, or test of spatial clustering, but to identify spatial clusters with their component points. Kintigh and Ammerman (see also Siegal and Roe 1986) have shown that this quantitative procedure identifies intuitively plausible clustering for both clear-cut artificial data and actual archaeological data.

In the simplest incarnation of pure locational clustering, a *k*-means non-hierarchical cluster analysis (Doran and Hodson 1975) is applied to the two-dimensional spatial coordinates of a set of points. The cluster analysis allocates each point into one of a specified number of clusters in a way that attempts to minimize a global goodness-of-fit measure, called sum-squared error (SSE). The SSE is simply the sum of the squared distances from each point to the center of the cluster to which it is assigned. The location of the center, or centroid, of a cluster is given by the mean for each dimension, e.g., east and north, over all points included in the cluster. Thus, for a given division of points into a given number of clusters, or *cluster configuration*, pure locational clustering locates the center of the cluster in space and identifies the points associated with that cluster.

Ordinarily, one makes an *a priori* determination of the maximum number of clusters that could conceivably be of interest and then obtains the cluster configurations for two clusters, three clusters, four clusters, and so forth up to the maximum number decided upon. Improved results, however, may often be obtained if a number of clusters somewhat larger than the maximum desired is requested. A plot of the SSE, expressed as a percentage of the SSE of the one-cluster solution, again the number of clusters (Figure 16) is used to identify the best, or "natural" clustering levels. The clustering levels of choice are those at which there are inflections in the SSE curve where the absolute value of the slope decreases. In Figure 16, for example, this is the case at three and eight clusters. The logic behind this criterion is that increasing the number of clusters beyond the number at which the inflection occurs does little to improve the goodness-of-fit, as measured by the reduction in the SSE, whereas fewer clusters have substantially higher SSE's.

Both the nature and the degree of spatial patterning can be assessed through a comparison of the SSE plot with the SSE plots of k-means analyses of randomized data (Kintigh and Ammerman 1982). Data that are significantly clustered at a given number of clusters should have an SSE value for that number of clusters below the SSE's obtained in analyses of randomized data. Points that are in fact randomly distributed will have SSE values within the range of SSE's for randomized data. Not so obviously, points that are evenly distributed will have SSE values greater than those obtained for randomized data.

Randomization is accomplished by creating a new data set with the same number of points wherein east (x) and north (y) coordinates are drawn separately and without replacement from the original data set. Randomized data sets thus derived have the same distribution of values on each spatial dimension, and hence have the same mean and standard deviation, and the same total, or one-cluster SSE as the original data set.

For the Mask Site, SSEs from *k*-means analyses of the actual and randomized data sets are presented in Figure 16. It can be seen that the actual distribution of points at the site is clustered when compared to the randomized data, and that inflections occur at the 3 and 8 cluster levels. The composition of the clusters at the 8 cluster level is tabulated in Table 7. The division of the 490 points into eight clusters is shown in Figure 17. In this figure, the cluster number associated with a point is displayed at its location, and each cluster centroid is plotted as a cross. A cluster "radius", which is simply the root of the mean of the squared distances from each point to its cluster center, and is denoted RMS, is drawn around each cluster.

Even this very simple pure locational clustering provides a rich body of relevant information that can figure directly in the interpretation of behavioral or formation processes that ultimately determined the distribution: the number of points in the clusters, the cluster sizes (RMS, in meters), locations (not listed in Table 7), and cluster composition relative to some reference variable, artifact class in this case.

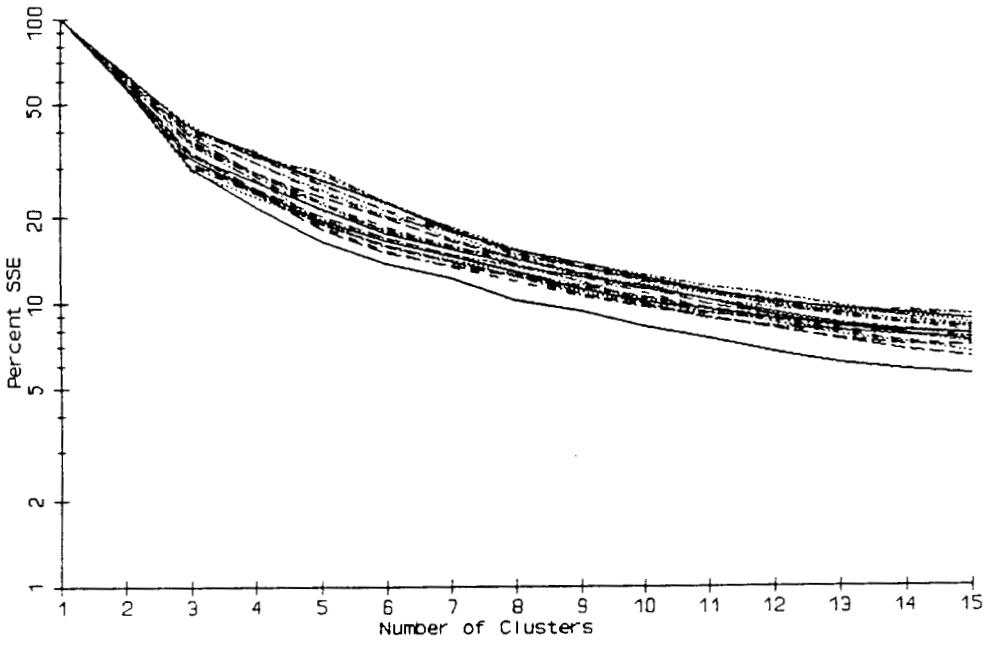

Figure 16. Mask Site (solid line) and randomized data (broken lines) pure locational clustering SSE plots.

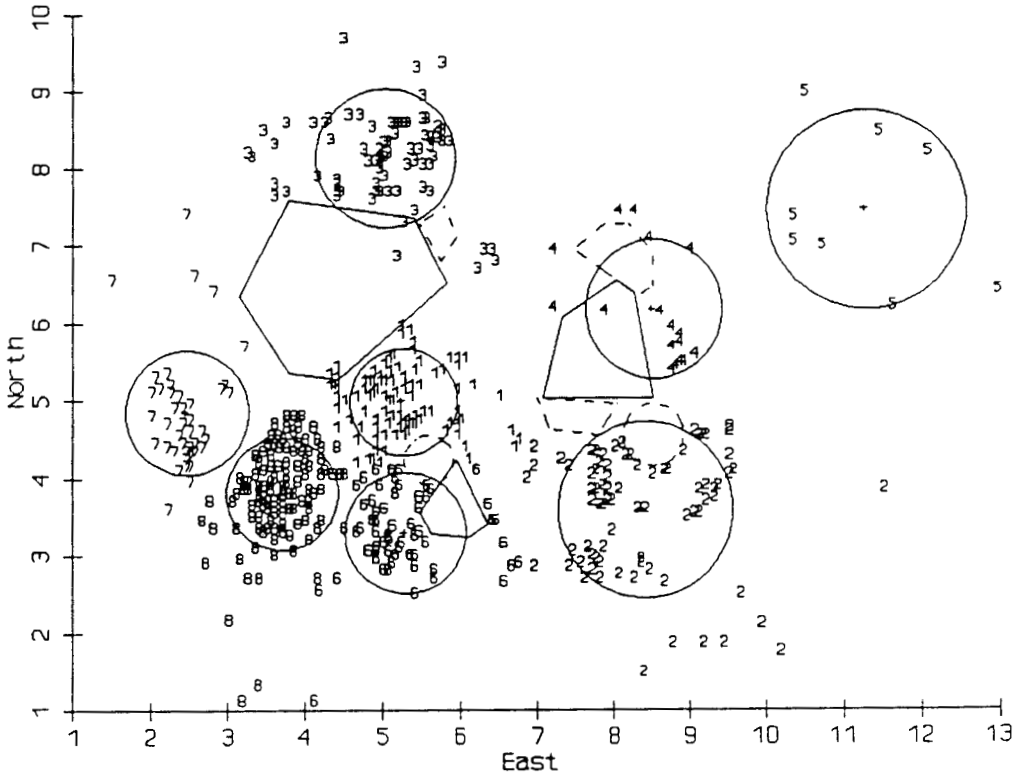

Figure 17. Mask Site pure locational clustering, plot of 8 cluster configuration.

It is apparent from Figure 17 and Table 7 that nearly all of the bone scrap is contained in four clusters (1, 2, 3, 6) immediately adjacent to the 3 major hearths identified by Binford. The large bone is mostly in clusters 2, 3, 7, 8, and is notably absent from the dense scatters of bone scrap in clusters 1 and 6. Cluster 5 in the northeast portion of the site has, relative to the other clusters, a radius that is large and an artifact density that is low, and is composed exclusively of tools. The wood scrap is found only in clusters 4 and 8, in the first case associated mainly with bone scrap, and in the second with projectiles. Nearly all the projectiles are in two adjacent clusters in the southwest portion of the site.

Comparison of the decline in SSE values for actual and randomized data against number of clusters provides direct information about overall patterning at different spatial scales. If the distribution is shown to be clustered, hypotheses related to formation processes or activity areas can be examined through an inspection of the location, size, point density, and composition of the individual clusters.

Pure locational clustering by artifact class

Pure spatial clustering, as described above, has provided the foundation for a variety of related analyses. As with some techniques discussed earlier, one can perform the analysis separately for each artifact class in order to examine the individual artifact classes for spatial patterning (Simek and Larick 1983, Simek 1984, Simek *et al.* 1985, Ammerman *et al.* 1985, Simek 1987). To the extent that one is interested in the distribution of artifacts in a particular class, this is an obvious and sensible approach.

However, one commonly wishes to know to what extent individual artifact classes display similar patterning, a question not directly answered by a pure locational clustering of all artifact classes combined. If pure locational clustering is performed on the artifact classes independently, we end up needing a way to assess the correspondence of pure locational clustering results. Methods of comparing or interpreting pure locational clustering configurations derived from separate analyses of artifact types remain largely intuitive and visual. Simek and Larick (1983) plot on a single map the cluster configurations with a fixed number of clusters derived from independent analyses of the artifact classes. Segregation of centroids and/or areas of overlap of the RMS radius circles of the different artifact classes are then interpreted.

However, when artifact classes are analyzed independently, the different classes may not be "naturally" clustered at the same number of clusters. For example, at the Mask Site , wood clearly has two clusters while projectiles clearly have three basic clusters (Figure 5). This deficiency, which was recognized by Simek and Larick in their 1983 paper, can

TABLE 7. PURE LOCATIONAL CLUSTERING OF THE MASK SITE; CLUSTER COMPOSITION, 8 CLUSTER SOLUTION. PERCENTAGES SHOWN ARE ROW PERCENTAGES, THE PERCENTAGE EACH ARTIFACT CLASS COMPRISES OF THE POINTS IN THE CLUSTER.

Cluster	N	RMS	Tool	Projectile	Wood Scrap	Large Bone	Bone Scrap
1	96	0.69	12/ 13%	0/ 0%	0/ 0%	1/ 1%	83/ 86%
2	79	1.14	8/ 10%	5/ 6%	0/ 0%	6/ 8%	60/ 76%
3	76	0.90	9/ 12%	0/ 0%	0/ 0%	8/ 11%	59/ 78%
4	17	0.90	1/ 6%	0/ 0%	7/ 41%	2/ 12%	7/ 41%
5	8	1.29	8/100%	0/ 0%	0/ 0%	0/ 0%	0/ 0%
6	60	0.78	9/ 15%	1/ 2%	0/ 0%	2/ 3%	48/ 80%
7	43	0.80	2/ 5%	31/ 72%	0/ 0%	5/ 12%	5/ 12%
8	111	0.73	3/ 3%	45/ 41%	50/ 45%	4/ 4%	9/ 8%

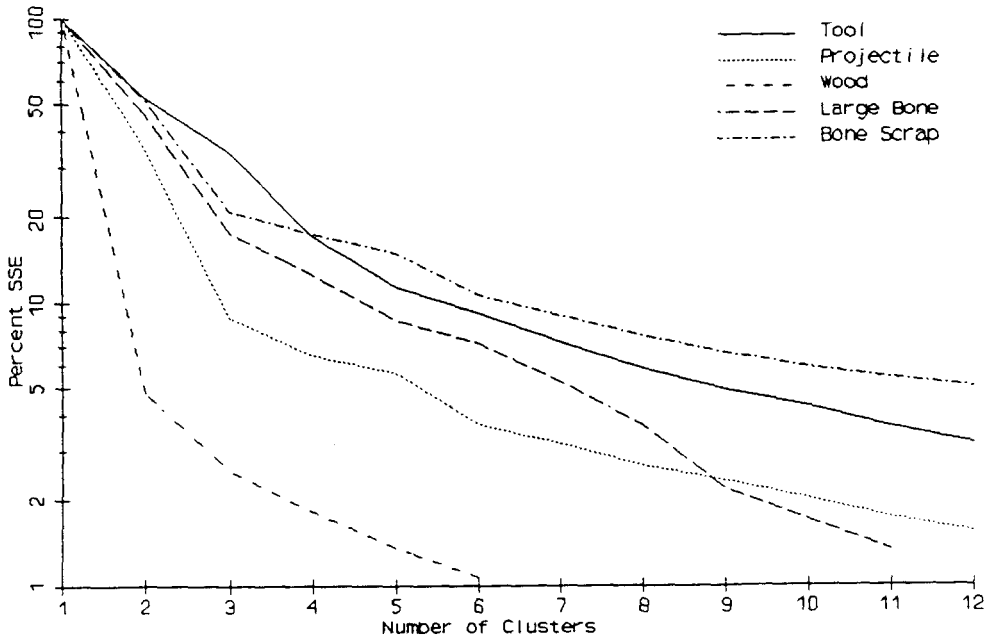

Figure 18. Mask Site pure locational clustering by type, SSE plot.

be dealt with using the natural breakpoints seen in the SSE curves to define clustering at generalized spatial scales, such as low, medium, and high resolution (Ammerman *et al.* 1985, Simek 1987). Thus, for the Mask Site, a consideration of low resolution clustering might include configurations of wood at 2 clusters, projectiles, large bone, and bone scrap at 3 clusters, and tools at 4 or 5 clusters (Figure 18).

Simek (1984, 1987:46-48, Ammerman *et al.* 1985) proposes a higher level concept of "zones" that are subjectively defined by the overlap of the RMS radius circles for different artifact class cluster configurations at a given level of spatial resolution. The zones are then used as the foci of interpretive discussions.

Homogeneity, spatial resolution, and depositional models

Simek (1984) made an important step in linking analytical methods with theoretical concepts of site formation and function based on Binford's (1978) and Yellen's (1977) ethnoarchaeological observations. He proposed models of deposition identified by differing patterns of compositional homogeneity of the zones at different levels of spatial resolution. In his procedure, artifact classes are cluster-analyzed separately and cluster zones are identified at differing levels of spatial resolution. The zones are then ranked, separately for each artifact class, according to the abundance of that artifact class in that zone. Finally, Simek measures homogeneity by the percentage of significant rank order correlations of all between-class correlations.

Koetje (1987:29-30, 43-44) has pointed out a serious defect with Simek's measure of homogeneity and replaces it with an alternative analytical procedure and refined depositional models. Rather than separately analyze the artifact classes, Koetje clusters all artifact classes together, thus eliminating the subjective steps involved in the identification of cluster zones at different spatial scales. He uses an index of cluster homogeneity that is a close numerical relative of Simpson's measure of diversity, adjusted by means of Monte

Carlo methods for what he calls the allocative effect—the mechanical relationships among the sample sizes of the artifact classes, the numbers of clusters in the original population, and his absolute measure of homogeneity.

For the Mask Site data, at the 8 cluster level the observed homogeneity, K′, is 0.52. Using Koetje's method, this may be compared with the distribution of K′ values computed for random allocations of the observed numbers of artifacts of each type to 8 clusters with sizes determined by the original analysis. Out of 1000 trials, the minimum K′ obtained by a random allocation of artifacts to clusters was 0.77 and the average was 0.78, indicating that the observed similarity in composition of the clusters to each other, or homogeneity, is much lower than would be expected if, with the same total counts per type, types were randomly assigned to artifact locations. This suggests non-random concentrations of artifacts of particular classes or sets of classes.

Identifying clustering levels to analyze

As discussed above, the standard method (Doran and Hodson 1975) of identifying the clustering levels to analyze has been to identify the points at which the SSE plot markedly levels out. While this semi-subjective method appears satisfactory for many purposes, some problems remain. One problem in identifying critical clustering levels appears when a SSE plot shows a gradual decline with no point of marked inflection.

An alternative approach, suggested by Gregg *et al.* (in press) may deal with data sets with no distinct inflections as well as with data sets exhibiting strong inflections. They suggest that for each clustering level, one should plot the difference between the percent SSE of the actual data against the mean percent SSE of the random runs for that number of clusters. Then, clustering levels with the largest differences between the actual and randomized data are selected.

While I believe that there is some merit to this approach, three issues not dealt with in Gregg *et al.* (in press) need to be considered. First, they used only two random runs in order to calculate an average random percent SSE. In general, this will be far too few random runs to obtain stable differences between actual and average random SSEs. For example, the SSE plot shown in Figure 16 shows the results of 25 random runs. It is easy to see that averaging an arbitrary 2 or 3 of those random runs could give rather different results than one sees with 25 random runs.

To use this strategy, one needs to do a sufficiently large number of random runs that random variation among these runs does not swamp the differences between the actual and random SSEs. An intuitive feel for the issue can be obtained by examining the spread of SSE values that result from doing many random runs. It can be evaluated more rigorously by looking at the standard deviation of the mean of the SSE obtained from the random runs at each clustering level, which is computed as the standard deviation of the sample of SSEs from the random runs divided by the square root of the number of random runs. The difference plot shown in Figure 19 shows the difference between the mean random SSE and the original data SSE and also plots this difference for plus and minus one standard deviation of the mean. It should be kept in mind, however, that experimentation has shown that a curve with a clear-cut peak does not always result from this computation.

Pure locational clustering, concluding remarks

Pure locational clustering is a powerful technique for the analysis of spatial data. It is untroubled by boundary problems and identifies the locations, sizes, composition, and members of the clusters. Simek, Koetje, and Larick have shown it to be subject to many

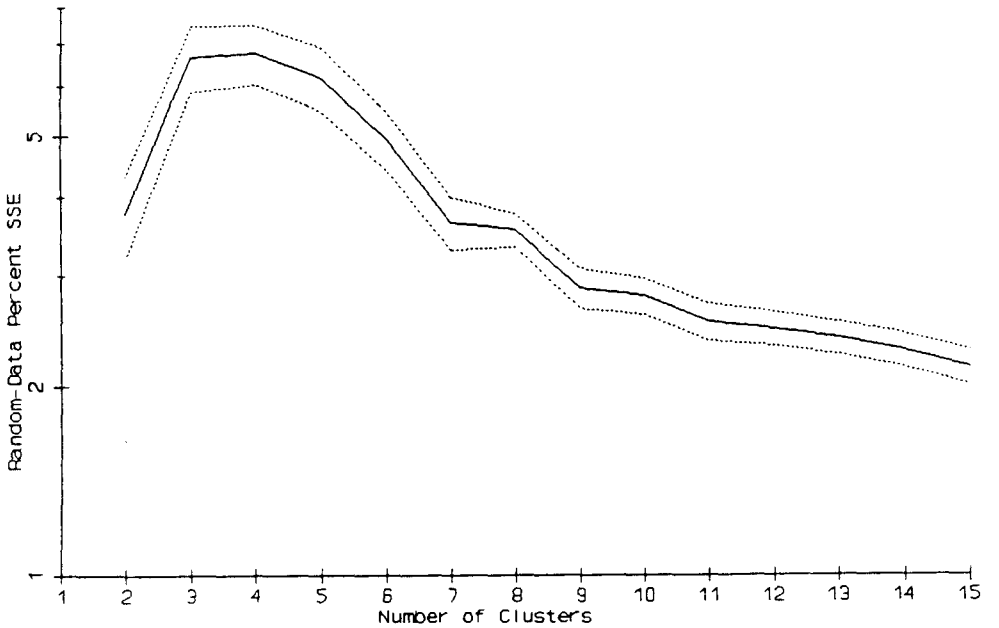

Figure 19. Mask Site pure locational clustering, random run-data SSE difference plot.

productive extensions. However, it does not provide a unique solution, a single set of artifact clusters, rather it identifies clustering at several spatial scales. This is advantageous in that there are often sound ethnographic reasons to believe that clustering may be evident at more than one spatial scale. It is problematic in that we sometimes do not have altogether satisfactory ways to identify the best levels of clustering to analyze, although it is not clear to me to what extent this is a problem with the method and to what extent the problem inheres in our conceptual models and in the data themselves.

Perhaps the most serious problem with pure locational clustering is that minimizing the global sum-squared error, the *k*-means algorithm attempts to form circular clusters. Although in some cases it performs satisfactorily on data sets with linear, oval, and crescentic clusters, and centers large clusters on them, or subdivides them into several clusters, its internal model is, at a fundamental level, inconsistent with some of our conceptual models. This is not a call to reject the method, but to keep its limitations in mind.

UNCONSTRAINED CLUSTERING

Whallon's (1984) unconstrained clustering is a spatial analysis technique that groups *areas* of a site in terms of their proportional artifact class composition. It is intended to be free of constraints with respect to cluster size, shape, and density.

As it is described by Whallon (1984, Gregg *et al.* in press), the first step in unconstrained clustering is to transform the artifact distribution map to one in which each vertex of a regular grid across the site is assigned a vector of smoothed densities of the artifact classes around that point. This is accomplished by centering a circular template (whose radius is the grid interval) at each grid vertex and calculating the absolute density of points of each type within the template. Because each template, in part, overlaps the templates applied to the 8 adjacent vertices, the densities are smoothed. From this grid of

smoothed densities, a contour map of artifact densities can be plotted, although this is not essential to unconstrained clustering.

Then, smoothed absolute densities are calculated for each original artifact location by interpolation, weighted by inverse squared distance, from the nearest grid points. In the next step, the vector of absolute densities at each artifact location is converted to a vector of proportional densities, i.e., the sum of the proportional, or relative, densities at each point is set to 1.0.

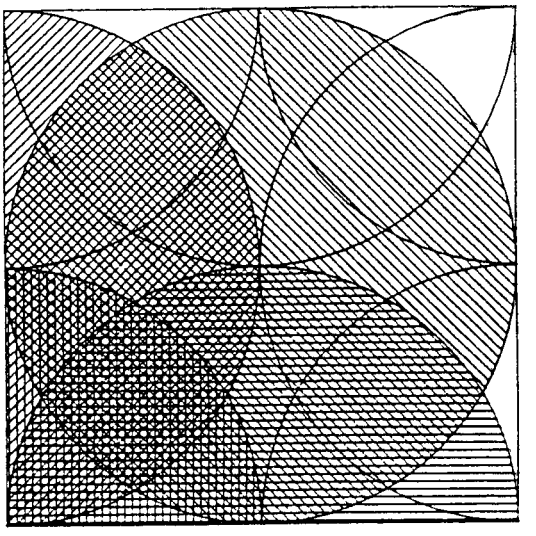

Figure 20. Unconstrained clustering circular template overlap.

Finally these vectors, representing relative densities of the artifact classes in the neighborhood of each artifact are subjected to a cluster analysis. The method used in Whallon (1984) is Ward's hierarchical clustering method. Based on the behavior of a clustering criterion, one or more clustering levels is selected for examination and the cluster number associated with each artifact location is plotted on a map. Spatial groupings of points belonging to the same clusters can then be identified by inspection and interpreted using descriptive statistics of the cluster analysis variables, e.g., the means and standard deviations of the relative densities of the artifact classes, calculated separately for the points included in each cluster. Whallon (1984) presents an unconstrained cluster analysis of the Mask Site data using both point-provenience and grid-count data and shows that these analyses lead to generally appropriate interpretations of the known ethnoarchaeological record.

As Whallon makes clear, unconstrained clustering is more an analytical strategy rather than a specific method. Within this general strategy, a variety of appropriate analytical decisions might be made concerning how to smooth and interpolate the densities, whether to standardize the data, the method used to cluster the density vectors, the clustering levels to analyze, or even the use of absolute rather than relative densities. Whallon also notes that the approach can be extended to grid data, often with little loss of resolution. In that case, the aggregation by grid unit effects a spatial smoothing of the densities.

Moving templates

While I believe unconstrained clustering to be a productive strategy for spatial analysis, some technical details of consequence have not yet received adequate consideration. Whallon's moving circular template is a case in point. A circle *seems* a reasonable shape, and as the template moves from vertex to vertex, it does indeed smooth the data by including each artifact in the domain of more than one template. However, as Susan Gregg has pointed out, it does this in a biased fashion (Gregg *et al.* in press).

Figure 20 shows the portions of the nine templates that impinge on any two-by-two unit square. The four templates that impinge on the upper left one-by-one unit square are hatched different ways, showing that some areas are sampled by four templates, some are sampled by three, and other areas are sampled only twice. Thus, some artifacts will count twice as much as others in the density calculations. In sparse portions areas of a distribution, it seems likely that this effect may be sufficiently important to affect interpretations.

A painful way to evaluate this phenomenon would be to move the origin of the smoothing grid somewhat and repeat the analysis. Pursuing this general strategy, perhaps a better solution would be to use a 2 × 2 unit square template to derive a density for the grid vertex at its center; in this way each artifact would be counted exactly 4 times. As will be seen below, switching to an analysis based on grid units, either unsmoothed or smoothed, eliminates the circular template problem and has some other theoretical advantages.

Relative densities in areas of low absolute density

Low density areas present a serious problem for the calculations of relative densities. This is easiest to see for the grid-count variant of unconstrained clustering. In grid-based unconstrained clustering, counts of the artifact classes are accumulated for each grid unit. These counts are then converted to proportions, or equivalently, to percentages. If a total grid unit count is relatively large, say, 20 or more, this is probably a reasonable step, as each artifact accounts for 5% or less. However, if only one or a few items are represented in the square, a *single* artifact may account for 20%, 33%, 50%, or as much as 100% in the vector of percentages.

For the typical real-world case, in which both low and moderate to high absolute densities are found on a site, this may present a serious problem for the subsequent step of cluster-analyzing the vectors of proportions. Here, the Euclidean distance between vectors of proportions used by the cluster analysis will often not yield a satisfactory measure of similarity between vectors because very small variation in absolute numbers of artifacts will have very large effects on the analysis. If areas that have low total counts also have a variety of artifact classes, single artifacts within low density grid units will have an enormous impact on the assignment of squares to individual clusters, and a very small total number of dispersed artifacts may significantly influence the overall cluster configurations obtained.

This is just as serious a concern for the standard point-based technique where the effect is only disguised by the additional calculations involved in smoothing and interpolation. Ultimately, the interpolated relative densities for sparse areas are still based on small total artifact counts. In this case, the circular template problem is particularly important, in that an artifact that lies in an area that is sampled 4 times by the smoothing algorithm will effect the relative density vector much more than an artifact that may be extremely close but is only counted twice.

Two obvious methods can be used together to alleviate this problem. First, one can exclude vectors for the point locations or grid units that are based on counts below some threshold. Second, one can expand the grid unit size or scale of smoothing so that larger total counts are involved in the computation of the proportions. Clearly, this involves a three-way tradeoff among spatial resolution, which varies with the grid unit size or smoothing scale, comprehensiveness, the inclusion or exclusion of units or points from the analysis, and robustness of the results, which vary with the use of proportions based on small counts.

A local density-based variant of unconstrained clustering

As Whallon indicates in his original article, the objective of the elaborate process of the smoothing and interpolation of the densities is to derive an estimate of the "... local density of any type of item at any specific point ..." (1984:246). Such a local density estimate can be derived in a way that is more direct and, I believe, theoretically preferable to Whallon's proposed procedure. This estimate can easily be acquired as a by-product of local density analysis computations.

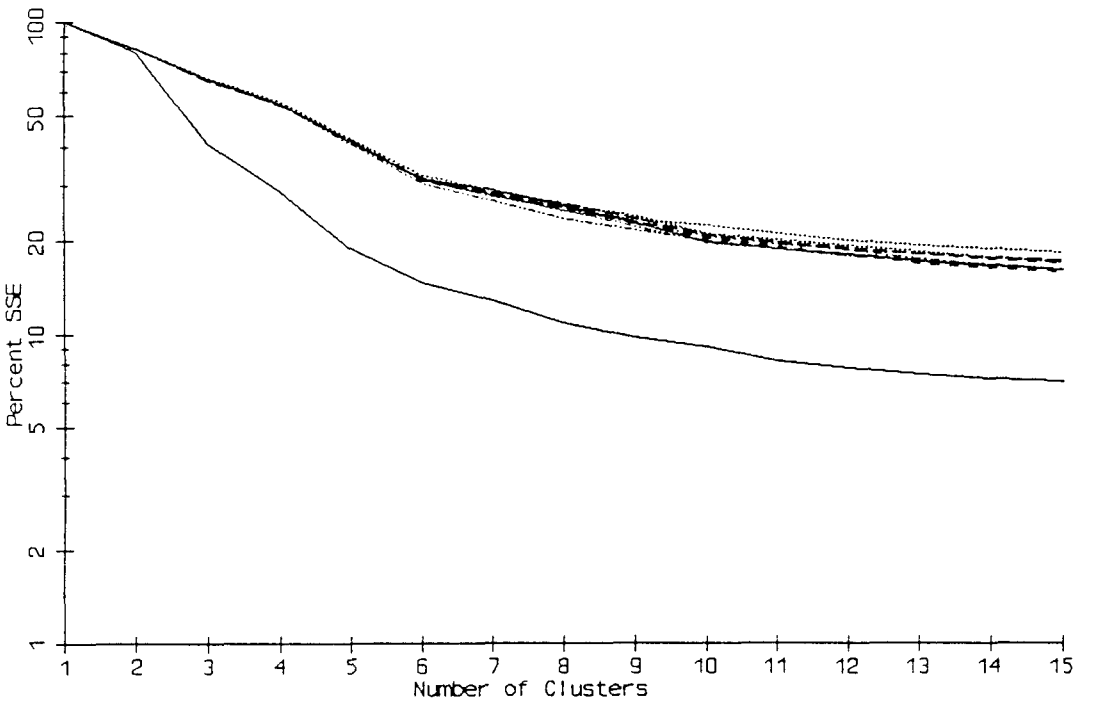

Figure 21. Mask Site unconstrained clustering using 1m radius local density, SSE plot.

The first step is to pick a radius for the neighborhood around each point on the basis of which the local densities are computed. Exactly as is done in local density analysis, each point is examined in turn and the count of points of each type within the fixed radius of the subject point is accumulated. Division of each count by the area of the circular neighborhood yields a vector of absolute density; division of each count by the total count of points in the circle yields a vector of relative densities.

This approach avoids the uneven sampling problem of the moving circular template, obviates the arbitrary aspect of grid placement, and effects a smoothing whose scale is perfectly clear-cut—the radius of the neighborhood. Points with a total count less than some threshold can easily be dropped from the analysis if desired.

This variant of unconstrained clustering was applied to the Mask Site, for a one meter radius neighborhood around all points, using a *k*-means clustering algorithm on standardized variables. The SSE plot (Figure 21) suggests examination of the 6 and 11 cluster levels. Results at the 6 cluster level are shown in Figure 22 and Table 8. It can be seen

TABLE 8. MEANS AND STANDARD DEVIATIONS FOR UNCONSTRAINED CLUSTERING OF THE MASK SITE WITH A 1M NEIGHBORHOOD RADIUS. THE PRIMARY COMPONENT OF EACH CLUSTER IN BOLDFACE; SECONDARY COMPONENTS IN ITALICS.

Cluster	N	Tool	Projectile	Wood Scrap	Large Bone	Bone Scrap
1	216	9/ 7%	1/ 3%	4/ 8%	2/ 2%	**84/ 10%**
2	3	0/ 0%	0/ 0%	0/ 0%	**89/ 16%**	*11/ 16%*
3	21	**93/ 12%**	0/ 0%	0/ 0%	5/ 8%	3/ 9%
4	54	2/ 2%	**75/ 10%**	8/ 10%	6/ 2%	9/ 6%
5	103	0/ 1%	*32/ 17%*	**46/ 9%**	3/ 5%	*18/ 13%*
6	93	6/ 6%	6/ 10%	0/ 1%	*12/ 5%*	**76/ 13%**

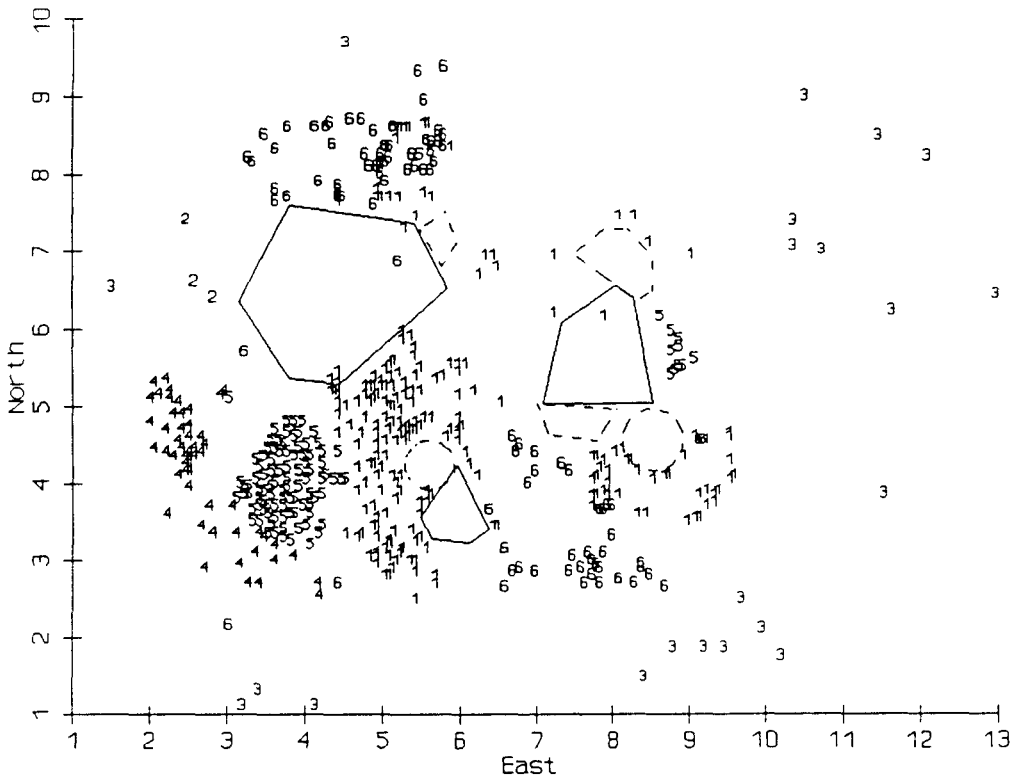

Figure 22. Mask Site unconstrained clustering using 1m radius local density, 6 cluster configuration plot.

that this analysis produces results quite similar to Whallon's seven cluster solution, in spite of the fact that the details of the analysis were quite different. An analysis in which neighborhoods with fewer than five points were dropped yielded similar results. However, this analysis makes clear that large bone is more strongly associated with the northeastern and southwestern hearth areas than the southeastern one.

According to Binford, bone scrap was dropped as individuals sat around the three major hearths. Points in cluster 1, dominated by bone scrap, surround each of the three major hearths. Larger bone elements were tossed over the shoulders of sitting individuals. The points in cluster 6, representing bone scrap associated with large bones, are associated with each major hearth, but are more distant than the cluster 1 points. The two activity-oriented clusters of wood scrap (cluster 5) and projectile components (cluster 4) are identified. Isolated tools (cluster 3) are on the periphery of the site, while three pieces of large bone are essentially isolated and form cluster 2.

Unsmoothed grid-based unconstrained clustering

An unsmoothed grid-based analysis of a 0.5m × 0.5m grid, with a minimum grid unit count of 1, using *k*-means clustering on standardized variables shows much the same patterning as the local-density analysis, with a few notable differences. The SSE plot as well as the random mean difference plot (Figure 23) clearly indicates that the 5 cluster solution should be examined (Figure 24, Table 9). In this plot the bone-scrap dominated cluster (1) still dominates the area around each major hearth; the projectile (4) and wood scrap (2) clusters are still clear; and the isolated tools (cluster 5) still show up at the edges of the site.

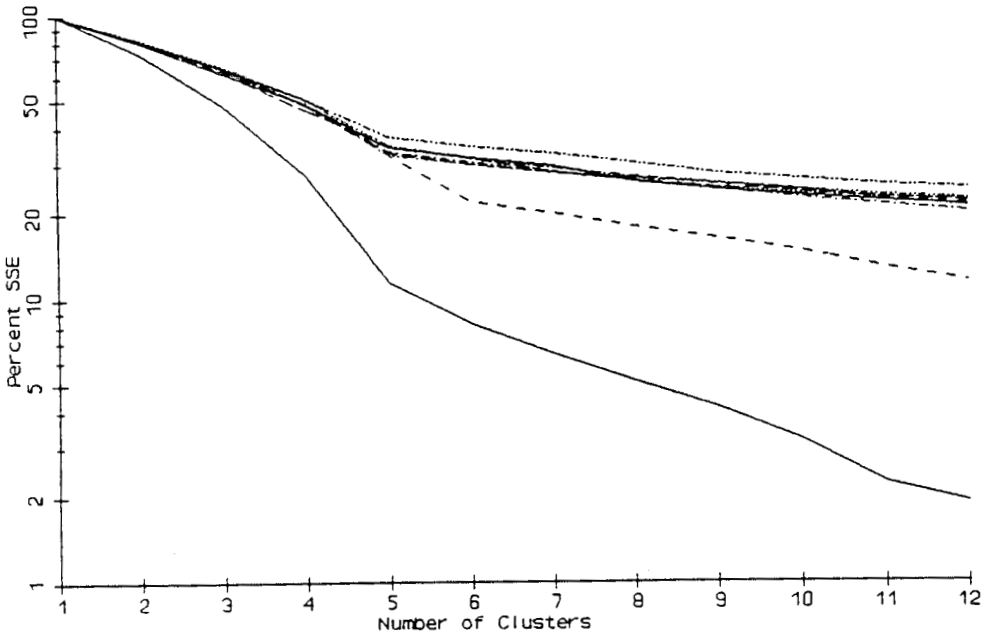

Figure 23. Mask Site unconstrained clustering using 1m unsmoothed grid, random run-data SSE difference plot.

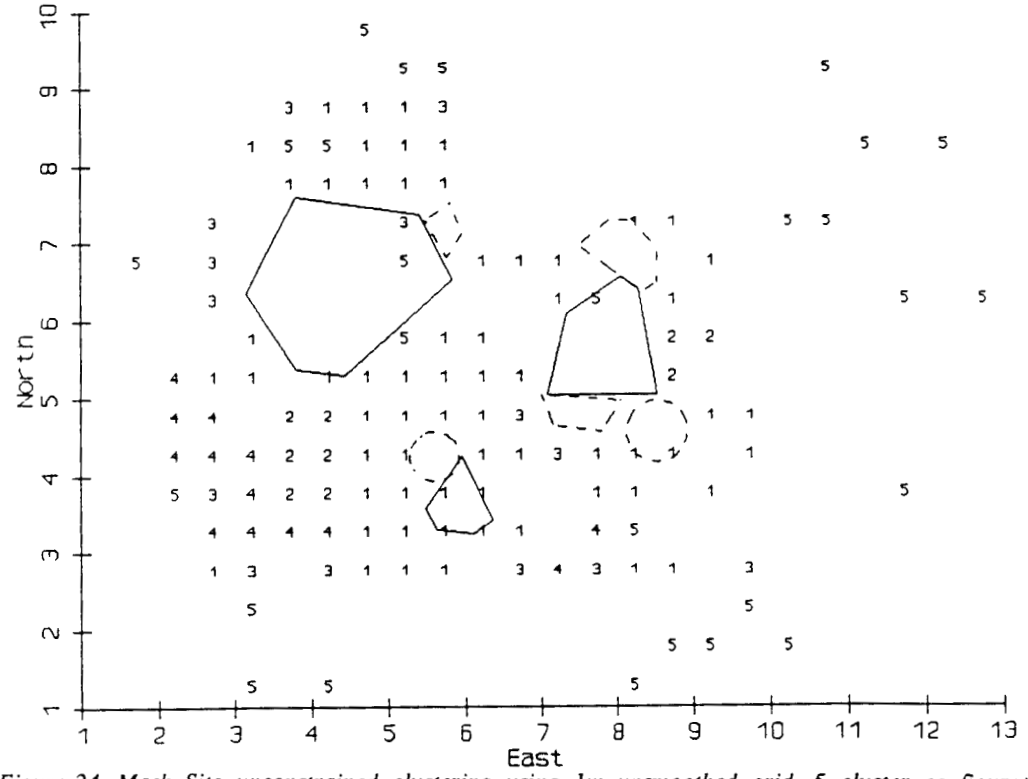

Figure 24. Mask Site unconstrained clustering using 1m unsmoothed grid, 5 cluster configuration plot.

However, in the grid-based analysis, it is more evident that the large bone (cluster 3) is scattered at some distance from the hearths than in the artifact-based analysis. Also in this analysis the clusters of tools in the interior of the site appear more clearly than they did in the point-based analyses. A grid-based analysis with a minimum grid count of five produced similar results.

Perhaps paradoxically, in this case, the grid-based approach yielded results more directly related to the observed behavior than artifact-based approaches at similar levels of clustering. It seems to me that there are some subtle but important advantages to an unsmoothed grid-based approach. Perhaps most important is that it removes the spatial autocorrelative effect introduced by smoothing, of whatever sort, in all point-based or smoothed-grid approaches.

In point-based approaches, the compositional vectors of points that are in the same neighborhood with respect to the smoothing procedure are *not* independent of one another. They are mathematically related in a way that guarantees that points that are relatively close together will generally end up in the same cluster. Thus, the patchiness of the clusters evident in plots like that shown in Figure 22 is in part a mechanical effect not an empirical result. This means that, for point-based approaches, only the appearance of *multiple* patches of a cluster in similar contexts but different areas of the site can be considered to be an clear sign of spatial patterning.

In contrast, when using an unsmoothed grid each artifact is counted only once, and the composition vectors associated with each grid unit are mathematically independent of each other. Thus, the appearance of adjacent grid units assigned to the same cluster indicates spatial patterning that cannot be a result of smoothing effects, and, as a consequence, leads to a stronger interpretation. In effect, use of an unsmoothed grid removes an important spatial constraint from artifact-based unconstrained clustering.

The descriptive statistics are also subject to a slightly cleaner interpretation in grid-based clustering of this sort. For example, from Table 9 one can conclude that of the 126 grid units ($31.5m^2$) that contain any artifacts, the 63 ($15.75m^2$) that form cluster 1 have an average of 95% bone scrap, 3% tools, and 2% large bone. An additional advantage is that the grid-based approach is usually much less computationally demanding than point-based approaches. Using *k*-means clustering, each run of the artifact-based approach took about 10 times the computation time of a single run of the half-meter grid analysis. A final procedural point about grid-based unconstrained clustering is that empty grid squares can be dropped from the analysis, saving computational energy with no loss of interpretive information.

Unconstrained clustering, concluding remarks

Unconstrained clustering, as a family of techniques, is a powerful tool for spatial analysis that is complementary to pure locational clustering (Gregg *et al.* in press).

TABLE 9. MEANS AND STANDARD DEVIATIONS FOR UNCONSTRAINED CLUSTERING OF THE MASK SITE WITH AN UNSMOOTHED, 0.5M × 0.5M GRID. PRIMARY COMPONENT OF EACH CLUSTER IN BOLDFACE; SECONDARY COMPONENTS IN ITALICS.

Cluster	N	Tool	Projectile	Wood Scrap	Large Bone	Bone Scrap
1	63	3/ 10%	0/ 0%	0/ 0%	2/ 7%	**95/ 12%**
2	9	0/ 0%	*10/ 18%*	**74/ 21%**	7/ 16%	9/ 13%
3	14	2/ 9%	*11/ 17%*	0/ 0%	**67/ 26%**	*20/ 24%*
4	13	0/ 0%	**84/ 19%**	2/ 7%	2/ 6%	*11/ 18%*
5	27	**100/ 0%**	0/ 0%	0/ 0%	0/ 0%	0/ 0%

Whereas pure locational clustering focuses on artifact density and location, unconstrained clustering focuses on class composition of areas of a site without regard to density.

Unconstrained clustering is well-suited for the examination of activity-based and many other behavioral/depositional models. As Whallon has pointed out, one outstanding problem in using this method—and probably all others—with activity area models is in recognizing overlap of the areas. I do not believe that Fourier-based methods (Graham 1980, Carr 1987) have been shown to deal adequately with this problem in practical cases, and see no unequivocal solution for it. However, when using unconstrained clustering, a zone with a partial overlap of activities can sometimes be seen in solutions with larger numbers of clusters by the division of the zones into "pure" and various gradations of "mixed" clusters as indicated by intermediate compositions in spatially adjacent clusters.

As with other methods, one must keep the limitations in mind. It is my sense that unsmoothed grid-based procedures have important advantages and that local density neighborhood densities are preferable to interpolated, smoothed densities in artifact-oriented analyses. With unconstrained clustering, several analytical steps are involved and many decisions with important analytical consequences must be carefully thought out.

CONCLUSIONS

The state of the art of intrasite spatial analysis in archaeology represents a constant interplay among questions asked by archaeologists, the quality of available data, models of behavior and deposition, and accessible methods. In recent years there have been important advances in all these areas. Further, I believe that promising directions for innovative research seem relatively clear.

For some kinds of research, the collection of highly precise provenience data has been the norm for many years; in others, standards for the collection of provenience information are rising. For purposes of spatial analysis, current and future excavations are less likely to be troubled by the precision of the provenience data than they are to be impeded by limited spatial extent of excavations, dictated by the time-consuming nature of modern recovery methods.

Until a few years ago, the model employed by quantitative techniques for comparison with actual data was a simple random distribution of points, and a statistically significant non-randomness was often interpreted as behaviorally significant patterning. Recently, considerable attention has been paid to the development of models of behavior that result in the spatial distributions we are left to deal with. Both ethnoarchaeological (e.g., Binford 1978; DeBoer and Lathrap 1979; Kent 1984, 1987; Yellen 1977) and archaeological (e.g., Boone 1987; Fletcher 1984; Leroi-Gourhan and Brézillon 1972; Kroll and Isaac 1984) studies have contributed to the development of important behavioral models that can be employed in quantitative spatial analysis.

A great number of papers have addressed the methodological issues. In addition to those cited here, Hietala (1984*) and Carr (1984) provide lengthy reviews. Some of the more recent methods were developed explicitly to deal with the lack of fit between realistic models of behavior and quantitative methods (Kintigh and Ammerman 1982; Simek 1984; Whallon 1984; Carr 1985; Koetje 1987). In addition, I have attempted to show that methods like nearest-neighbor analysis can be adapted in such a way that they can be informative with respect to more contemporary concerns. I think there is reason to expect relatively rapid progress on methodological issues because as sophisticated spatial analysis techniques become more widely available on microcomputers, they become more widely used and better understood.

As with other concerns of quantitative archaeology, it seems to me that too often our methods have been allowed to dictate our questions rather than the other way around (Kintigh 1987). However, over the last several years, we have greatly refined our questions

and, as has been shown here, Monte Carlo techniques can be used with existing procedures to obtain answers to reformulated, and more relevant questions. We are now concerned with local patterning and multiple levels of patterning, as well as with global patterning. Assessing the fit between model and data remains a major topic. A healthy concern with the validation of results currently is being coupled with inductive, pattern searching techniques, which I believe to be essential for spatial analysis in archaeology. The problem of mixing or overlap of deposit types is perhaps the outstanding issue before us. There is much work to be done, but much progress has been made.

REFERENCES

Ammerman, A.J., K.W. Kintigh, and J.F. Simek 1985
Recent developments in the application of the *k*-means approach to spatial analysis. In: *Human Uses of Flint and Chert: Papers from the Fourth International Flint Symposium*, Volume 2. Cambridge University Press, Cambridge.

Binford, L.R. 1978
Dimensional analysis of behavior and site structure: learning from an Eskimo hunting stand. *American Antiquity* 43:30-61.

Boone, J.L. 1987
Defining and measuring midden catchment. *American Antiquity* 52:336-345.

Carr, C. 1984
The nature of organization of intrasite archaeological records and spatial analytic approaches to their investigation. In: M.B. Schiffer (ed.), *Advances in Archaeological Method and Theory, Volume 7*, p. 103-222. Academic Press, New York.

Carr, C. 1985
Alternative models, alternative techniques: variable approaches to intrasite spatial analysis. In: C. Carr (ed.), *For Concordance in Archaeological Analysis*, p. 302-473. Westport Press, Kansas City.

Carr, C. 1987
Dissecting intrasite palimpsests using Fourier methods. In: S. Kent (ed.), *Method and Theory for Activity Area Research: An Ethnoarchaeological Approach*, p. 236-291. Columbia University Press, New York.

Clark, P.J., and F.C. Evans 1954
Distance to nearest neighbor as a measure of spatial relationships in populations. *Ecology* 35:445-453.

Clark, G.A. 1979
Spatial association at Liencres, an early Holocene open site on the Santander coast, north-central Spain. *Computer Graphics and Archaeology. Arizona State University Anthropological Research Papers* 15:121-143.

Clarke, D.L. (ed.) 1977
Spatial Archaeology. Academic Press, New York.

Cowgill, G.L., J.H. Altschul, and R.S. Sload 1984
Spatial analysis of Teotihuacán: a Mesoamerican metropolis. In: H.J. Hietala (ed.), *Intrasite Spatial Analysis in Archaeology*, p. 154-195. Cambridge University Press, Cambridge.

DeBoer, W.R., and D.W. Lathrap 1979
The making and breaking of Shipibo-Conibo ceramics. In: C. Kramer (ed.), *Ethnoarchaeology: Implications of Ethnography for Archaeology*, p. 102-138. Columbia University Press, New York.

Doran, J.E., and F.R. Hodson 1975
Mathematics and Computers in Archaeology. Harvard University Press, Cambridge.

Fletcher, R. 1984
Identifying spatial disorder: a case study of a Mongol fort. In: H.J. Hietala (ed.), *Intrasite Spatial Analysis in Archaeology*, p. 196-223. Cambridge University Press, Cambridge.

Graham, I. 1980
Spectral analysis and distance methods in the study of archaeological distributions. *Journal of Archaeological Science* 7:105-129.

Gregg, S., K.W. Kintigh, and R. Whallon in press
Linking ethnoarchaeological interpretation and archaeological data: the sensitivity of spatial analytical methods to post-depositional disturbance. In: E.M. Kroll and T.D. Price (eds.), *The Interpretation of Archaeological Spatial Patterning*. Plenum, New York.

Hanson, G., and A. Goodyear 1985
The shared-tool method of spatial analysis: applications at the Brand Site. Manuscript on file with the Department of Anthropology, Arizona State University.

Hietala, H.J. (ed.) 1984[a]
Intrasite Spatial Analysis in Archaeology. Cambridge University Press, Cambridge.

Hietala, H.J. 1984[b]
Variations on a categorical data theme: local and global considerations with Near-Eastern Paleolithic applications. In: H.J. Hietala (ed.), *Intrasite Spatial Analysis in Archaeology*, p. 44-53. Cambridge University Press, Cambridge.

Hivernel, F., and I. Hodder 1984
Analysis of artifact distribution at Ngenyn (Kenya): depositional and post depositional effects. In: H.J. Hietala (ed.), *Intrasite Spatial Analysis in Archaeology*, p. 97-115. Cambridge University Press, Cambridge.

Hodder, I., and E. Okell 1978
A new method for assessing the association between distributions of points in archaeology. In: I. Hodder (ed.), *Simulation Studies in Archaeology*, p. 97-107. Cambridge University Press, Cambridge.

Hodder, I., and C. Orton. 1976
Spatial Analysis in Archaeology. Cambridge University Press, Cambridge.

Johnson, I. 1976
Contribution Méthodologique à l'Étude de la Répartition des Vestiges dans des Niveaux Archéologiques. Thesis for obtaining a Diplome des Études Superieurs, Université de Bordeaux I.

Johnson, I. 1984
Cell frequency recording and analysis of artifact distributions. In: H.J. Hietala (ed.), *Intrasite Spatial Analysis in Archaeology*, p. 75-96. Cambridge University Press, Cambridge.

Kent, S. 1984
Analyzing Activity Areas: An Ethnoarchaeological Study of the Use of Space. University of New Mexico Press, Albuquerque.

Kent, S. (ed.). 1987
Method and Theory for Activity Area Research: An Ethnoarchaeological Approach. Columbia University Press, New York.

Kintigh, K.W. 1987
Quantitative methods designed for archaeological problems. In: M.S. Aldenderfer (ed.), *Quantitative Research in Archaeology: Progress and Prospects*, p. 126-134. Sage Publications, Newbury Park.

Kintigh, K.W., and A.J. Ammerman. 1982
Heuristic approaches to spatial analysis in archaeology. *American Antiquity* 47:31-63.

Koetje, T.A. 1987
Spatial Patterns in Magdalenian Open Air Site from Isle Valley, Southwestern France. British Archaeological Reports, International Series, 346, Oxford.

Kroll, E.M., and G.Ll. Isaac 1984
Configurations of artifacts and bones at early Pleistocene sites in East Africa. In: H.J. Hietala (ed.), *Intrasite Spatial Analysis in Archaeology*, p. 4-31. Cambridge University Press, Cambridge.

Leroi-Gourhan, A., and M. Brézillon 1972
Fouilles de Pincevent: essai d'analyse ethnographique d'un habitat Magdalénien (la section 36). *Gallia Préhistoire*, Supplement 7. C.N.R.S., Paris.

Pinder, D., I. Shimada, and D. Gregory 1979
The nearest-neighbor statistic: archaeological application and new developments. *American Antiquity* 44:430-445.

Siegal, P.E., and P.G. Roe 1985
Shipibo archaeo-ethnography: site formation processes and archaeological interpretation. *World Archaeology* 18:96-115.

Simek, J.F. 1984
A K-means Approach to the Analysis of Spatial Structure in Upper Paleolithic Habitation Sites: Le Flageolet I and Pincevent Section 36. British Archaeological Reports, International Series, 205, Oxford.

Simek, J.F. 1987
Spatial order and behavioural change in the French Paleolithic. *Antiquity* 61:25-40.

Simek, J.F., A.J. Ammerman, and K.W. Kintigh 1985
Explorations in heuristic spatial analysis: analyzing the structure of material accumulations over space. In: A. Voorrips and S.H. Loving (eds.), *To Pattern the Past*, p. 229-247. PACT 11, Strasbourg.

Simek, J.F., and R.R. Larick 1983
The recognition of multiple spatial patterns: a case study from the French Paleolithic. *Journal of Archaeological Science* 10:165-180.

Whallon, R. 1974
Spatial analysis of occupation floors II: the application of nearest-neighbor analysis. *American Antiquity* 39:16-34.

Whallon, R. 1984
Unconstrained clustering for the analysis of spatial distributions in archaeology. In: H.J. Hietala (ed.), *Intrasite Spatial Analysis in Archaeology*, p. 242-277. Cambridge University Press, Cambridge.

Whallon, R. 1987
Simple statistics. In: M.S. Aldenderfer (ed.), *Quantitative Research in Archaeology: Progress and Prospects*, p. 135-150. Sage Publications, Newbury Park.

Yellen, J.E. 1977
Archaeological Approaches to the Present: Models for Reconstructing the Past. Academic Press, New York.

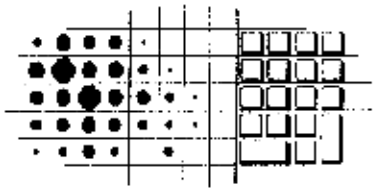

Studies in Modern Archaeology

Vol. 3

Mathematics and Information Science In Archaeology:

A Flexible Framework

Albertus Voorrips (ed.)

HOLOS

Bonn 1990