Nearest neighbor methods

Philip M. Dixon

Nearest neighbor methods

Nearest neighbor (NN) methods include at least six different groups of statistical methods. All have in common the idea that some aspect of the similarity between a point and its NN can be used to make useful inferences. In some cases, the similarity is the distance between the point and its NN; in others, the appropriate similarity is based on other identifying characteristics of the points. I will discuss in detail NN methods for spatial point processes and field experiments because these are commonly used in biology and environmetrics. I will very briefly discuss NN designs for field experiments, in which each pair of treatments occurs as neighbors equally frequently. I will not discuss NN estimates of probability density functions (pdfs) [23], NN methods for discrimination or **classification** [59], or NN linkage (i.e. simple linkage) in hierarchical clustering [32, pp. 57–60]. Although these last three methods have been applied to environmetric data, they are much more general.

Nearest Neighbor Methods for Spatial Point Processes

Spatial point process data describe the locations of 'interesting' events (*see* **Point processes, spatial**) and (possibly) some information about each event. Some examples include locations of tree trunks [52], locations of bird nests [11], locations of pottery shards, and locations of cancer cases [20]. I will focus on the most common case where the location is recorded in two dimensions (x, y). Similar techniques can be used for three-dimensional data (e.g. locations of galaxies in space) or one-dimensional data (e.g. nesting sites along a coastline or along a riverbank). Usually, the locations of all events in a defined area are observed (completely mapped data), but occasionally only a subset of locations is observed (sparsely sampled data). Univariate point process data include only the locations of the events; marked point process data include additional information about the event at each location [65]. For example, the species may be recorded for each tree, some cultural identification may be recorded for each pottery shard, and nest success or nest failure may be recorded for each bird nest.

Location or marked location data can be used to answer many different sorts of questions. The scientific context for a question depends on the area of application, but the questions can be grouped into general categories. One very common category of question concerns the spatial pattern of the observations. Are the locations spatially clustered? Do they tend to be regularly distributed, or are they random (i.e a realization of a homogeneous **Poisson process**)? A second common set of questions concerns the relationships between different types of events in a marked point process. Do two different species of tree tend to occur together? Are locations of cancer cases more clustered than a random subset of a control group (*see* **Disease mapping**)? A third set of questions deals with the density (number of events per unit area). What is the average density of trees in an area? What does a map of density look like? Methods to answer each of these types of question are discussed in the following sections.

Theoretical treatments of NN methods for spatial point patterns can be found in [18], [25], [58] and [65]. Applications of NN methods can be found in many articles and books, including [40], [53], [64] and [66].

Describing and Testing Spatial Patterns Using Completely Mapped Data

Describing and testing spatial patterns of locations has a long history. Historically, the primary concern was with the question of randomness [1, 2, 15, 48]. Are locations randomly distributed throughout the study area (i.e. are the locations a realization of a Poisson process with homogeneous intensity), or do the locations indicate some structure (i.e. clustering or repulsion between locations)? Because of the many connotations of randomness and the importance of a homogeneous Poisson process as a benchmark, it is commonly called *complete spatial randomness* (CSR).

In this section I will describe NN tests based on completely mapped data. Locations of all events are recorded in an arbitrary study region. Often, the study region is square or rectangular, but this is not a requirement. Tests for the less common case of sampled data are described in the next section.

2   Nearest neighbor methods

Tests Based on Mean Nearest Neighbor Distance

The distances between NNs provide information about the pattern of points. Define W as the distance from a randomly chosen event to the nearest other event in a homogeneous Poisson process with intensity (expected number of points per unit area; *see* **Poisson intensity**) of ρ. The pdf and cumulative **distribution function** (cdf) of W are

$$g(w) = 2\rho\pi w \exp(-\rho\pi w^2) \tag{1}$$

$$G(w) = 1 - \exp(-\rho\pi w^2) \tag{2}$$

so the mean and the variance of W are

$$E(W) = \frac{1}{2\rho^{1/2}} \tag{3}$$

and

$$\text{var}(W) = \frac{4-\pi}{4\pi\rho} \tag{4}$$

Based on these moments, Clark and Evans [15] proposed a test of CSR. The conditional moments, $E(W|\hat{\rho})$ and $\text{var}(W|\hat{\rho})$, are calculated by substituting the observed density, $\hat{\rho}$ (total number of points divided by the total area of the study region), into (3) and (4). The observed mean NN distance, $\overline{w}$, is computed by identifying the NN of each point, finding the distance between NNs, then averaging. Clark and Evans proposed that the standardized mean

$$Z_{\text{CE}} = \frac{\overline{w} - E[W|\hat{\rho}]}{(N^{-1}\text{var}[W|\hat{\rho}])^{1/2}} \tag{5}$$

has a standard normal distribution if the process is a homogeneous Poisson process [15].

The Z_{CE} statistic and the many users of it ignore two problems: nonindependence of some NN distances, and **edge effects**. In a completely mapped area, many of the distances between NNs are correlated. The problem is most severe with reflexive NNs. Two points, A and B, are reflexive NNs when B is the NN of A, and A is the NN of B [16]. Other authors have called these 'isolated NNs' [51] or 'mutual NNs' [61]. When A and B are reflexive NNs, each point has the value of W, which inflates the variance of the mean NN distance. This problem is not restricted to a few points. When points exhibit CSR in two dimensions, approximately 62.15% of the points are reflexive NNs [16].

Edge effects arise because the distribution of W (2) assumes an unbounded area, but the observed NN distances are calculated from points in a defined study area. When a point is near the edge of the study area, it is possible that the true NN is a point just outside the study area, not a more distant point that happens to be in the study area. Edge effects lead to overestimation (positive bias) of the mean NN distance. Edge effects can be practically important; neglecting them can alter conclusions about the spatial pattern [10].

Edge effects may be minimized by including a buffer area that surrounds the primary study area [15]. NN distances are calculated only for points in the primary study area, but locations in the buffer area are available as potential NNs. With a sufficiently large buffer area, this approach can eliminate edge effects, but it is wasteful since an appropriately large buffer area may contain many locations.

Using simulations, Donnelly [31] derived edge-corrected approximations to $E(W|\hat{\rho})$ and $\text{var}(\overline{w}|\hat{\rho})$ when the study region is rectangular:

$$E(W|\hat{\rho}) \approx 0.5 \left(\frac{A}{N}\right)^{1/2}$$

$$+ 0.0514\frac{P}{N} + 0.041\frac{P^{3/2}}{N} \tag{6}$$

$$\text{var}(\overline{w}|\hat{\rho}) \approx 0.0703\frac{A}{N^2} + 0.037\frac{PA^{1/2}}{N^{5/2}} \tag{7}$$

where N is the observed number of points, A is the area of the study region, and P is the total perimeter of the study region. These approximations can be used to test CSR by substituting them into (5), then comparing the resulting z**-statistic** to a standard normal distribution. This test has reasonable power to detect departures from CSR [57].

One difficulty with tests based on the mean NN distance is that the mean is just a single summary of the pattern. Two point patterns may have the same mean NN distance, but one exhibits CSR and the other does not. One such pattern would have a few patches of clustered points and an appropriate number of widely scattered individuals. The points in the clusters have small NN distances, but the widely scattered individuals have large NN distances. With the appropriate mix of clustered and scattered points, the mean NN distance could be exactly that given by (3).

Distribution of Nearest Neighbor Distances

An alternative is to consider the entire distribution, $G(w)$, of NN distance [24]. CSR can be tested by comparing the observed distribution function of NN distances, $\widehat{G}(w)$, to the theoretical cdf (2). A variety of test statistics have been suggested, including Kolmogorov–Smirnov type statistics: $\sup_w |\widehat{G}(w) - G(w)|$, Cramer–von Mises type statistics: $\int_w [\widehat{G}(w) - G(w)]^2$, or Anderson–Darling type statistics: $\int_w [\widehat{G}(w) - G(w)]^2 / [G(w)(1 - G(w))]$. The Kolmogorov–Smirnov statistic seems to be the most commonly used. The usual critical values for the one-sample Kolmogorov–Smirnov test are not appropriate here because of nonindependence of NN distances, especially for reflexive NNs, as discussed above. Instead, a Monte Carlo test must be used [28].

The Monte Carlo approach computes the α-level critical value or the ***P* value** by simulation (*see* **Simulation and Monte Carlo methods**). The number of points, N, is fixed at the observed number. N random locations are simulated as a realization of CSR in the study area, and the Kolmogorov–Smirnov (or other) test statistic is computed. This is repeated R times to give R values from the sampling distribution of the test statistic under the null hypothesis of CSR. The one-sided α-level critical value is the $\alpha(R + 1)$th largest value from the sampling distribution. The P value is computed as $(1 +$ number of random values larger than the observed value$)/(1 + R)$ [5].

Edge effects complicate the estimation of $\widehat{G}(w)$. One solution is to include a border strip, as discussed above, but this may ignore a considerable amount of information. A variety of edge-corrected estimators of $G(w)$ has been proposed; four of them are summarized in [18, pp. 613–614, 637–638]. The edge-corrected cdf can also be estimated by a Kaplan–Meier estimator [3]. Edge corrections reduce the bias in the estimator, but they increase the sampling variance [36].

Although edge-corrected estimators are needed if the observed distribution function is to be compared with the theoretical distribution function under CSR (2) they may not be needed for a Monte Carlo test of CSR. A more powerful test of CSR is to use a nonedge-corrected estimator (2) and compare the biased estimate of $\widehat{G}(w)$ to the biased mean, $\overline{G}(w)$, under CSR [36]. The biased mean $\overline{G}(w)$ and pointwise prediction intervals are computed by

simulation. Values of $\widehat{G}(w)$ above the simulated mean indicate clustering of points (an excess of short distances to NNs). Values of $\widehat{G}(w)$ below the simulated mean indicate regularity (few to no points with short distances to NNs).

Graphical diagnostics based on the empirical distribution of NN distances provide additional information about the spatial process. The most common graphical diagnostic is a quantile–quantile plot of either $G(w)$ or $\overline{G}(w)$ on the x-axis and $\widehat{G}(w)$ on the y-axis. $G(w)$ would be used when the object is to evaluate the fit of the theoretical NN distribution; $\overline{G}(w)$ would be used when the theoretical NN distribution is unknown and the object is to evaluate the fit to a process that can be simulated.

Monte Carlo simulation of the spatial process provides both an estimate of $\overline{G}(w)$ and the sampling distribution of $\widehat{G}(w)$ conditional on a specific spatial process. **Quantiles** of the sampling distribution of $\widehat{G}(w)$ can be calculated for interesting distances, w. If the spatial process is simulated many times (e.g. 199 or 999), the quantiles can be estimated relatively precisely. The 0.05 and 0.95 quantiles can be approximated by the minimum and maximum of $R = 19$ simulations. The 0.01 and 0.99 quantiles can be approximated by the minimum and maximum of $R = 99$ simulations.

The simulated mean and quantiles of $\widehat{G}(w)$ are plotted against the observed $\widehat{G}(w)$. If the observed curve of $\widehat{G}(w)$ falls entirely between the lower and upper bounds, then there is no evidence against the null hypothesis (e.g. that the locations are a realization of CSR). An excursion outside the bounds indicates a departure from CSR. If the observed $\widehat{G}(w)$ falls below the lower bound at short distances, there are too few NNs at short distances, which is consistent with a regular pattern, or one where there is inhibition of nearby points. If the observed $\widehat{G}(w)$ lies above the upper bound at short distances, there are too many NNs at short distances, which is consistent with a clustered process.

The Monte Carlo approach is not limited to testing CSR. It can be used to evaluate the fit of any process that can be simulated. For example, a set of locations might be compared with a **Poisson cluster process** or a Strauss process [25, 31].

NN methods have been extended in a variety of ways. Tests can be based on other functions of the NN distances (e.g. squared [12] or smallest [62] NN distances), but such tests have not been widely used.

Distances to second NN, or the third NN, or perhaps an even further neighbor have been suggested as a way to look at patterns on a larger scale. The set of cdfs of distance to the NN, distance to the second NN, . . . , distance to nth NN is related to **Ripley's K function**, another commonly used method to analyze spatial patterns [56; 65, p. 267]. Finally, the distance between a randomly chosen point and the nearest event also provides information about the spatial pattern [25, 66].

Point–Event Distances

The point–event distribution, $F(x)$, considers the distance between a randomly chosen location (not the location of an event) and the nearest event. This can be estimated by choosing m locations in the study area and computing the distance from each location to its NN. As with $G(w)$, edge effects complicate estimation of the cdf. An edge-corrected estimator is

$$\hat{F}_{\mathrm{R}}(x) = \frac{\text{number of } (x_{\mathrm{i}} \leq x, d_{\mathrm{i}} > x)}{\text{number of } (d_{\mathrm{i}} > x)} \qquad (8)$$

where x_{i} is the distance between a point and its neighboring event, and d_{i} is the distance between a point and its nearest boundary. When the events are a realization of CSR, X, the point–event distance, and W, the NN distance, have the same distribution, so $F(x) = 1 - \exp(-\rho\pi x^2)$. However, the effects of deviations from randomness on $F(x)$ are opposite to those on $G(w)$. Values of $\hat{F}(x)$ above the expected value indicate regularity. Values below the expected value indicate clustering.

The NN distance distribution, point–event distance distribution, and Ripley's K function provide different insights into the spatial pattern. The NN distribution function, $\hat{G}(w)$, is slightly more powerful at detecting departures from CSR in the direction of regularity [24]. The point–event distribution function provides information about the empty space between points. It appears to be the more powerful method for detecting departures in the direction of clustering [25, 66]. Ripley's K function simultaneously examines the spatial pattern at many distance scales and is now the most popular approach for completely mapped data. However, it is possible to construct point processes with the same $G(w)$ but a different Ripley's $K(t)$ function [65, p. 267]. Conversely, Baddeley and Silverman [4] illustrate two processes with the same Ripley's K function but very different

NN distance distributions and point–event distance distributions.

Are Events Points or Circles?

All the distributional theory of the previous sections assumes that events occupy no space. Treating events as points assumes that it is possible for two events (e.g. locations of tree trunks or bird nests) to be an infinitesimally small distance from each other. The point assumption is reasonable when the area of the events is small relative to the spacing between the events. The assumption is likely to be appropriate for bird nests (generally small) or tree trunks (generally low density), but not for ant nests (large size relative to the density of nests in an area). If events are incorrectly assumed to be points, the analysis of the spatial pattern indicates a tendency to regularity because two events do not occur within a small distance (the physical size of the event) of each other.

An approximation to the mean event–event distance, $\overline{W}$, for nonoverlapping circles under CSR is

$$\overline{W} \approx \frac{d + \exp(-\rho\pi d^2)[1 - \Phi(2\rho\pi d)^{1/2}]}{\sqrt{\rho}} \qquad (9)$$

where d and ρ are, respectively, the diameter of the circles and the number of circles per unit area [63]. In (9) a low intensity, ρ, and a small and constant diameter, d, are assumed. The distribution, $G(w)$, of NN distances for nonoverlapping circles can be estimated by Monte Carlo simulation by using a sequential inhibition algorithm [24]. The distribution of event–event distances can be complicated when the circles are large or the density is high. For example, it may not be possible to fit the required number of large circles into the study area.

Algorithms and Computing

The simplest way to compute NN distances is by direct enumeration; that is, computing the distances between all pairs of points, then reporting the smallest distance for each point as the NN distance. For a large number of points, this becomes impractical, and more efficient algorithms have been developed to approach the problem. Possible approaches include subdividing the region into smaller subregions [47], computing the Direchlet **tessellation** (also known as the Voronoi tessellation, Thiessen tessellation, and other less common names [66, p. 96]) and using that

to identify NNs, or computing the quadtree, a sorted matrix of locations that simplifies the search for NNs, and using that to identify NNs [35]. Murtagh [47] reviews the properties of these algorithms.

Functions or procedures for NN spatial analysis are included in few statistical programs, but direct enumeration is very simple to program when needed. Packages of functions for spatial point pattern analysis usually include functions for point–point and point–event analyses. Many of these are S or **S-PLUS** libraries, for example Splancs [60], spatial [67], and **S + SPATIALSTATS** [45].

Example: Trees in a Swamp Hardwood Forest

Figure 1 shows the locations of all 630 trees (stems > 11.5 cm diameter at breast height) and the locations of 91 cypress trees in a 1 ha plot of swamp hardwood forest in South Carolina, USA. There are 13 different tree species in this plot, but over 75% of the stems are one of three species: black gum (*Nyssa sylvatica*), water tupelo (*Nyssa aquatica*), or bald cypress (*Taxodium distichum*). Visually, trees seem to be scattered randomly throughout the plot, but cypress trees seem to be clustered in three bands. NN statistics provide a way to test the hypothesis that stems are randomly distributed throughout the plot. I will illustrate tests based on mean NN distance, $G(w)$, and $F(x)$ using the locations of all 630 trees and the locations of the 91 cypress trees.

For all 630 tree locations, the mean NN distance is 1.99 m. If 630 points were randomly distributed in a 200 m × 50 m rectangle, the expected NN distance is 2.034 m, with a **standard error** (SE) of 0.044 m, using Donnelly's approximations, (6) and (7). There is no evidence of departure from CSR ($z = -0.973$ with a two-sided P value of 0.33). The effect of the edge corrections is minimal probably because the plot is large and the NN distance is small. The uncorrected expected NN distance is 1.992 m, with an SE of 0.042 m, using (3) and (4).

Conclusions using the distribution of NN distances, $\widehat{G}(w)$, are similar. $\widehat{G}(w)$ was estimated without edge corrections, so $\widehat{G}(w)$ must be compared with simulated values, not the theoretical expectation (2). The observed cdf, the theoretical expected value (2), and the average simulated cdf are very similar (Figure 2a), although the observed $\widehat{G}(w)$ is slightly larger than the expected value at short distances. The differences can be seen more clearly
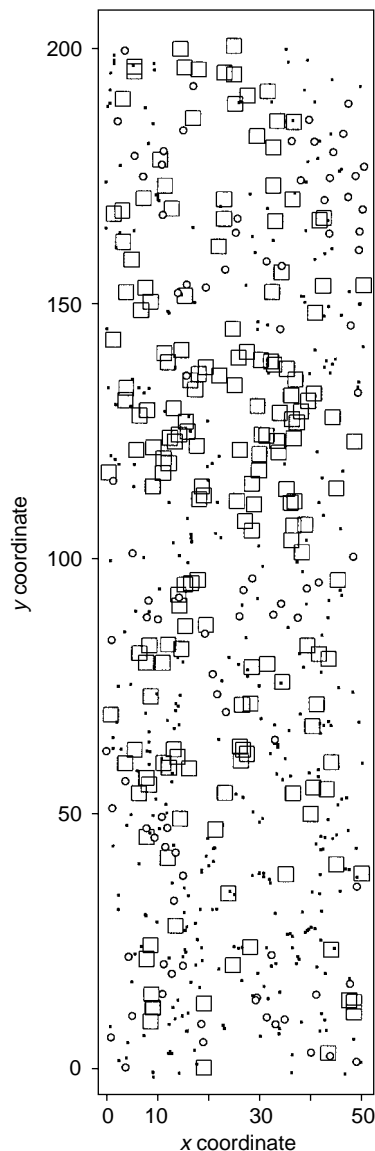

Figure 1 Marked plot of tree locations in a 50 m × 200 m plot of hardwood swamp in South Carolina, USA. Circles are locations of cypress trees, squares are locations of black gum trees, and dots are locations of any other species

if $1 - \exp(-\rho\pi w^2)$ (2) is subtracted from all curves (Figure 2b). Although $\widehat{G}(w)$ is larger than both the theoretical expected value and the average simulated value, it lies within the pointwise 95% confidence limits. None of the three summary statistics (Kolmogorov–Smirnov, Cramer–von Mises, or

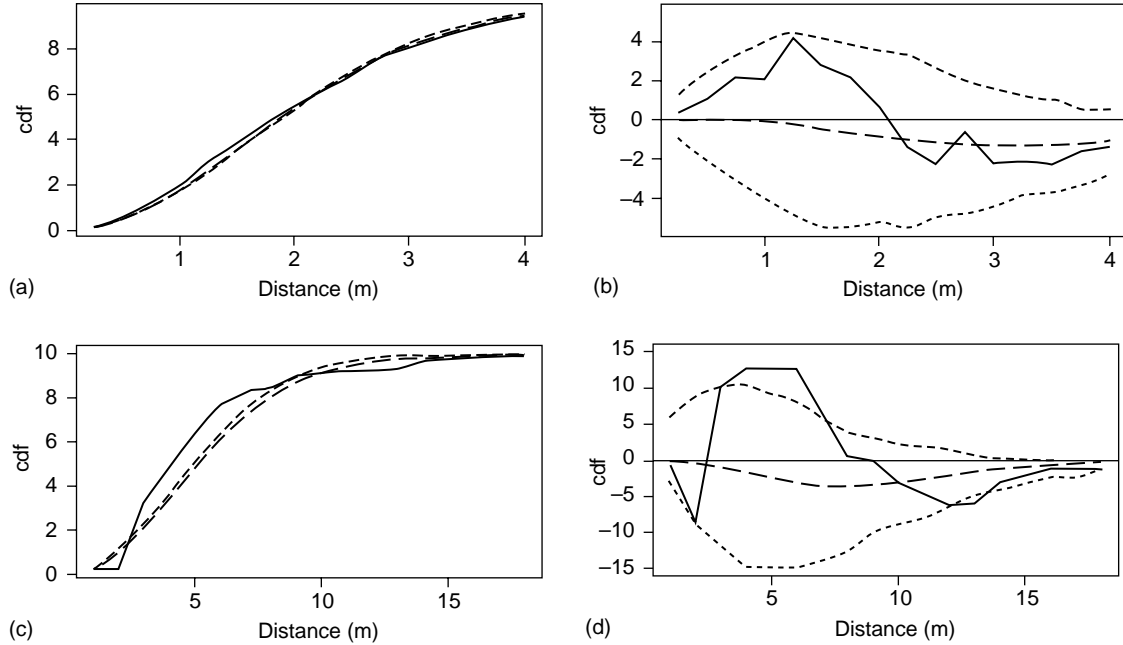**Figure 2** Cdf plots of NN distances for all trees and for cypress trees only, in a 50 m × 200 m plot (see Figure 1). (a) All trees: comparison of observed cdf (solid line), theoretical cdf (dotted line) and average simulated cdf under CSR (dashed line). (b) All trees: comparison of observed cdf (solid line) and average simulated cdf (dashed line). The dotted lines represent the pointwise 0.025 and 0.975 quantiles of the simulated cdf under CSR. For clarity, the theoretical cdf is subtracted from all plotted cdfs. (c) Cypress trees only: comparison of observed cdf (solid line), theoretical cdf (dotted line) and average simulated cdf under CSR (dashed line). (d) Cypress trees only: comparison of observed cdf (solid line) and average simulated cdf (dashed line). The dotted lines represent the pointwise 0.025 and 0.975 quantiles of the simulated cdf under CSR. For clarity, the theoretical cdf is subtracted from all plotted cdfs

Anderson–Darling) is significant at $\alpha = 0.05$. For example, the observed Kolmogorov–Smirnov test statistic of 0.044 is less than the simulated 90th percentile, 0.052. The estimated P values for the Kolmogorov–Smirnov, Cramer–von Mises, and Anderson–Darling test statistics are 0.19, 0.31, and 0.40, respectively.

In contrast, the distribution of point–event distances suggests there is some clustering of tree locations. The observed $\widehat{F}(x)$ is below the theoretical and average simulated curves (Figure 3a, b) and outside the pointwise 95% confidence bounds at large distances (Figure 3b). Because $\widehat{F}(x)$ falls below the expected values, distances from randomly chosen points are stochastically greater than expected if events exhibited CSR. The greater than expected abundance of large empty spaces provides evidence of clustering of the events. The P values for the Kolmogorov–Smirnov, Cramer–von Mises, and Anderson–Darling summary statistics range from

0.004 to 0.009. The conclusion of some evidence for clustering of all tree locations matches the conclusion using Ripley's K function.

For the 91 cypress trees, the mean NN distance is 5.08 m, which is slightly smaller than the edge-corrected expected distance of 5.55 m, with an SE of 0.33 m. Using the NN distance, there is no evidence of a nonrandom distribution; the z-statistic is −1.41, with a two-sided P value of 0.16. The effect of the edge corrections is larger when the density of points is smaller. The uncorrected expected NN distance for the 91 cypress trees is 5.24 m, with an SE of 0.29 m.

However, the distributions of $\widehat{G}(w)$ and $\widehat{F}(x)$ for the 91 cypress trees provide evidence of clustering of cypress trees. There is an unusually large number of NN distances between 3 m and 7 m (Figures 2c and 2d). This excess is significant; $\widehat{G}(w)$ is at or above the pointwise 0.975 quantiles of simulated values (Figure 2d). This excess is consistent with clustering of cypress trees. There are
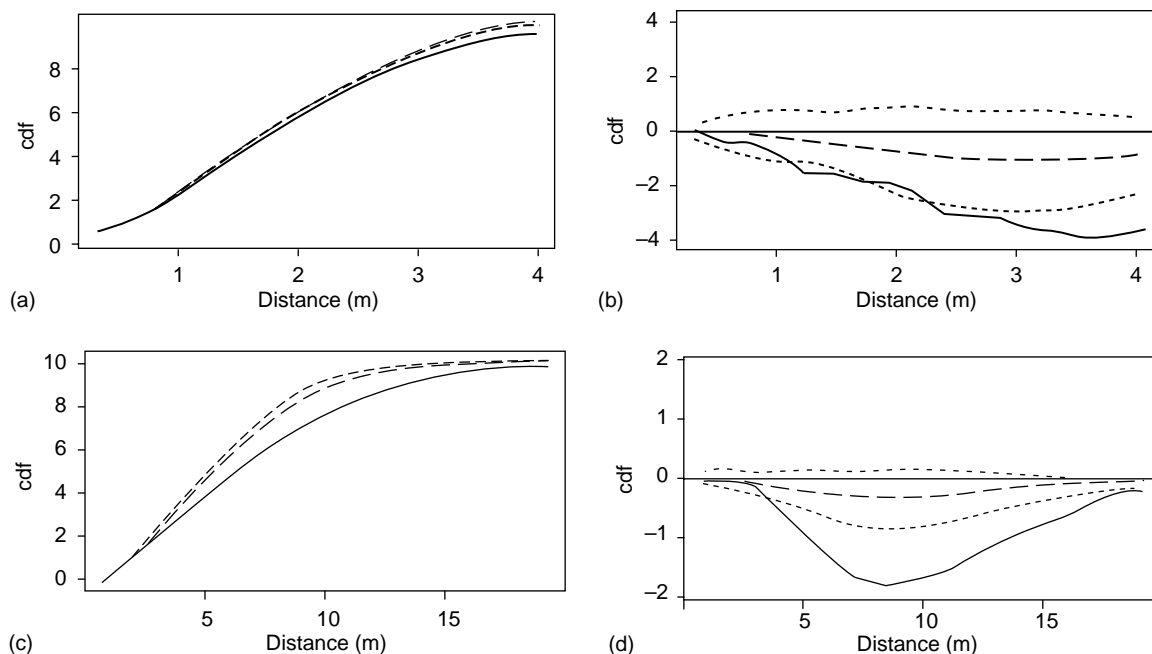
Figure 3   Cdf plots of point–event distances for all trees and for cypress trees only, in a 50 m × 200 m plot (see Figure 1). (a) All trees: comparison of observed cdf (solid line), theoretical cdf (dotted line) and average simulated cdf under CSR (dashed line). (b) All trees: comparison of observed cdf (solid line) and average simulated cdf (dashed line). The dotted lines represent the pointwise 0.025 and 0.975 quantiles of the simulated cdf under CSR. For clarity, the theoretical cdf is subtracted from all plotted cdfs. (c) Cypress trees only: comparison of the observed cdf (solid line), theoretical cdf (dotted line) and average simulated cdf under CSR (dashed line). (d) Cypress trees only: comparison of observed cdf (solid line) and average simulated cdf (dashed line). The dotted lines represent the pointwise 0.025 and 0.975 quantiles of the simulated cdf under CSR. For clarity, the theoretical cdf is subtracted from all plotted cdfs

also significant fewer (at least pointwise) NN distances at 13 m. All three summary statistics are significant (Kolmogorov–Smirnov P value $= 0.034$, Cramer–von Mises P value $= 0.007$, Anderson–Darling P value $= 0.011$). The point–event distances are stochastically greater than expected under CSR (Figures 3c and 3d). The observed distribution, $\widehat{F}(x)$, lies outside the pointwise 95% confidence bands for many distances. All three summary statistics are highly significant ($P = 0.001$). The conclusion that cypress trees are strongly clustered matches that using Ripley's K function.

Directed Tests

The tests in the previous section are general tests of CSR against an unspecified alternative. Other tests may be more powerful when the alternative is more specific (e.g. events are associated with specific sites,

or the density of events increases from east to west). Association between point events and a nonpoint stochastic process can be tested using the NN distance from each event to the nearest part of the second process [7]. However, most directed tests [43, 55, 68] use features other than NN distance.

Describing and Testing Spatial Patterns with Use of a Sample of Nearest Neighbor Distances

Although mapped data are often easy to collect, a statistician might view this as wasteful because of the high degree of correlation between NN distances (e.g. the perfect correlation for pairs of reflexive NNs). An alternative is to measure NN distances only on a random sample of individuals. Because the NN distances are calculated from a simple random sample of points,

the distributional theory for both the mean NN distance and the distribution function is much simpler. Many different tests of CSR have been developed for use with a random sample of point–event, NN, or point–event–event distances. These are summarized and evaluated in [66, pp. 59–64] and [18, pp. 602–614].

The most straightforward way to select a random sample of NN distances is to enumerate all individuals in the statistical population, select a simple random sample of events, and measure the distances from the selected events to their NNs. This can be time-consuming and is usually impractical [66]. Enumeration can be avoided by clever use of subregions (described by Byth and Ripley [14]), or by randomly selecting points (not events). The distance from the randomly selected point to the nearest event is a random sample from the distribution of point–event distances, $F(x)$, but the distance from that event to its closest event is not a random sample from $G(w)$ because the point–event and event–event distances are correlated. The distributions of all quantities in the point–event–event sample when events exhibit CSR have been derived [17].

An alternative that is easy to implement in the field is the T-square sample [8], illustrated in [18] and [66], a modified point–event–event sample. A point, A, is randomly chosen, and the NN, B, is found. Then, the study area is divided into two half planes by a line through B and perpendicular to AB (hence the name, T-square). Attention is restricted to the half plane that does not contain point A. The distance to the NN, Z, of B in that half plane is measured. When points exhibit CSR, Z and X (the distance from point A to NN B) are independent, and the distribution of $Z/\sqrt{2}$ is the same as the distribution of NN distances, $G(w)$ [8].

Estimating Density

A random sample of point to NN distances can be used to estimate ρ, the average density of events in the study area. When events exhibit CSR, the **maximum likelihood estimate** is

$$\hat{\rho} = n \left(\pi \sum_{i=1}^{n} X_i^2 \right)^{-1} \qquad (10)$$

where X_i is the distance from a randomly chosen point i to its NN [46]. An unbiased estimate is [54]:

$$\hat{\rho}_{\text{P}} = (n - 1) \left(\pi \sum_{i=1}^{n} X_i^2 \right)^{-1} \qquad (11)$$

Both estimators are very dependent on the CSR assumption and can be biased if locations are clustered or regularly distributed.

Many other estimators have been proposed. Upton and Fingleton [66, pp. 118–133] summarize and provide examples of calculations for many of these. Byth [13] evaluated the robustness of many estimators to deviations from CSR. She recommended an estimator, $\hat{\rho}_{\text{T}}$, based on two quantities from T-square sampling: X, the distance from point to nearest event, and Z, the distance to the NN in a half plane:

$$\hat{\rho}_{\text{T}} = n^2 \left[2\sqrt{2} \left(\sum_{i=1}^{n} X_i \right) \left(\sum_{i=1}^{n} Z_i \right) \right]^{-1} \qquad (12)$$

Nearest Neighbor Methods to Examine Spatial Patterns of More Than One Type of Point

Additional information about an event is often available. For example, tree locations might be marked with the species of tree (a mark with discrete levels) or the size of the tree (a mark with continuous levels). The methods in the previous sections can be used to analyze patterns in all events (ignoring the marks) or subgroups of events (e.g. just species A or just trees larger than 50 cm). However, other interesting questions could be asked about the relationship between the two (or more) sets of locations.

Multivariate spatial point patterns are those where events can be classified into different types, that is, the marks are discrete [18, p. 707]. Usually, the number of different types is small; bivariate patterns, with two types of marks, are the most common. Some questions that could be asked about point processes with discrete marks are as follows:

1. Are the processes that generate locations with different marks independent?
2. Are marks randomly assigned to locations? Conditional on the observed locations of superposition of the two marked processes, are the marks independent?

3. Are marks segregated? Are locations with one type of mark surrounded by locations with the same mark?

These questions about the relationships between processes make no assumptions about the marginal pattern of each process. In particular, either process (or the superposition of processes) may be independent, clustered, or regularly distributed. The two general methods to answer these questions are the comparison of distribution functions [39, 44] and the NN contingency table [30, 52]. Other approaches include Ripley's K functions or parametric point process models.

Define the following multitype extensions of the point–event and NN distances. X_i is the distance from a randomly chosen point to the nearest event with mark i, with cdf $F_i(x)$. W_{ij} is the distance from an event with mark i to the nearest event with mark j, with cdf $G_{ij}(w)$. If the process with mark i is independent of the process with mark j, then

$$F_i(x) = G_{ji}(x) \qquad (13)$$

$$F_j(x) = G_{ij}(x) \qquad (14)$$

and X_i and X_j are independent [27, 39]. Note that property (13) does not imply property (14), so two tests are needed [39].

For sparsely sampled data, Goodall [39] suggests a t test of $\overline{X}_i = \overline{X}_j$. Diggle and Cox [27] consider nonparametric versions of the t test, tests of equality of distribution, and tests of the correlation between X_i and X_j. Details and a comparison of the tests are given in [27]. Analyses of completely mapped data tend to focus on the comparison of distribution functions in (13) and (14). Monte Carlo tests are used because of the nonindependence of point–event and event–event distances.

Two different simulation methods could be used in the Monte Carlo test. The choice depends on the null hypothesis. If the null hypothesis is independence between marks (question 1, above), then toroidal shifts or some parametric model should be used to generate the randomization distribution. If the null hypothesis is random assignment of marks conditional on the set of events, then random labeling of events should be used to generate the randomization distribution. In general, these two hypotheses are not equivalent and the sampling distributions are not the same.

Nearest Neighbor Contingency Tables

The NN contingency table focuses on the ecologically important question of segregation [30, 52]. This table describes marks of events and their NNs, not the distance between them (Table 1). In sparsely sampled data, the counts (N_{AA}, N_{AB}, N_{BA}, and N_{BB}) are independent Poisson random variables or conditionally independent given the row marginal totals (N_A and N_B) under the null hypothesis of random labeling [52]. The hypothesis can be tested with a traditional 1 degree of freedom (df) χ^2 test of independence [52] (*see* **Categorical data**).

In completely mapped data, the sampling distribution of the counts is different [16]. If events are randomly labeled, the expected values of the counts depend only on the number of each type of event (N_A and N_B) and the total number of events, N (Table 2) [30]. The variances and covariances depend on the number of events of each type, the number of reflexive NNs, and the number of shared NNs; for a derivation of the formula, see [30]. The first two moments of the cell counts can be used to test for segregation of type A events [$N_{AA} > E(N_{AA})$], test for segregation of type B events [$N_{BB} > E(N_{BB})$],

Table 1   Cell counts in an NN contingency table

Mark of point	Mark of neighbor		Total
	A	B	
A	N_{AA}	N_{AB}	N_A
B	N_{BA}	N_{BB}	N_B
Total	M_A	M_B	N

N is the number of points in the spatial pattern. N_A and N_B are the number of type A and type B points. N_{AA} is the number of type A points with type A NNs. N_{AB} is the number of type A points with type B NNs. N_{BA} and N_{BB} are the number of type B points with type A or type B NNs, respectively. M_A and M_B are the column sums, i.e. the total number of points with type A neighbors or type B neighbors, respectively

Table 2   Expected cell counts for NN contingency table for completely mapped data

From:	To:	
	A	B
A	$\dfrac{N_A(N_A - 1)}{N - 1}$	$\dfrac{N_A N_B}{N - 1}$
B	$\dfrac{N_B N_A}{N - 1}$	$\dfrac{N_B(N_B - 1)}{N - 1}$

See Table 1 for definitions of M, N_A and N_B

or construct an omnibus 2 df χ^2 test of random labeling. If the numbers of points are large, then the distributions of test statistics can be adequately approximated by asymptotic normal and χ^2 distributions. If the number of points is small, then the distributions should be determined by Monte Carlo simulation.

Patterns with k marks ($k > 2$) can be analyzed by considering all pairs of marks two at a time (using distance methods or 2×2 NN contingency tables), or by considering the $k \times k$ contingency table. The expected counts and their variances under random labeling follow the same form as those for a 2×2 contingency table, but there are more possible forms for the covariance between two counts. Details are given in [29].

Other approaches that have been suggested for the analysis of multitype point processes include the comparison of bivariate Ripley's K functions [26, 44], empty space methods [44] (comparisons of point–event distance distributions), and mark correlation functions [65].

Example, Part 2

Cypress and black gum are two of the three abundant species in the 1-ha plot of swamp forest considered in Example 1. An interesting ecological question is whether these two species are spatially segregated; that is, do cypress trees tend to be found near other cypress trees and do black gum trees tend to be found near other black gum trees? The marked plot of locations (Figure 1) suggests that cypress trees and black gum occur in different clusters. Confirming this requires an analysis of the bivariate spatial pattern. The three tests that will be illustrated are the comparisons of cdfs (13–14) [27], the independence of distances [27], and the NN contingency table [30]. The ecological background suggests that random labeling is the more appropriate null hypothesis.

The cdf of distances from randomly chosen points to the nearest black gum tree, $F_G(x)$, and the cdf of point–event distances to the nearest cypress tree, $F_C(x)$, were estimated without edge corrections by using a randomly located grid of points [14]. The cdfs of distances from black gum trees to the nearest cypress tree, $G_{GC}(x)$, and from cypress trees to the nearest black gum tree, $G_{CG}(x)$, were also estimated without edge corrections. Both species show the same
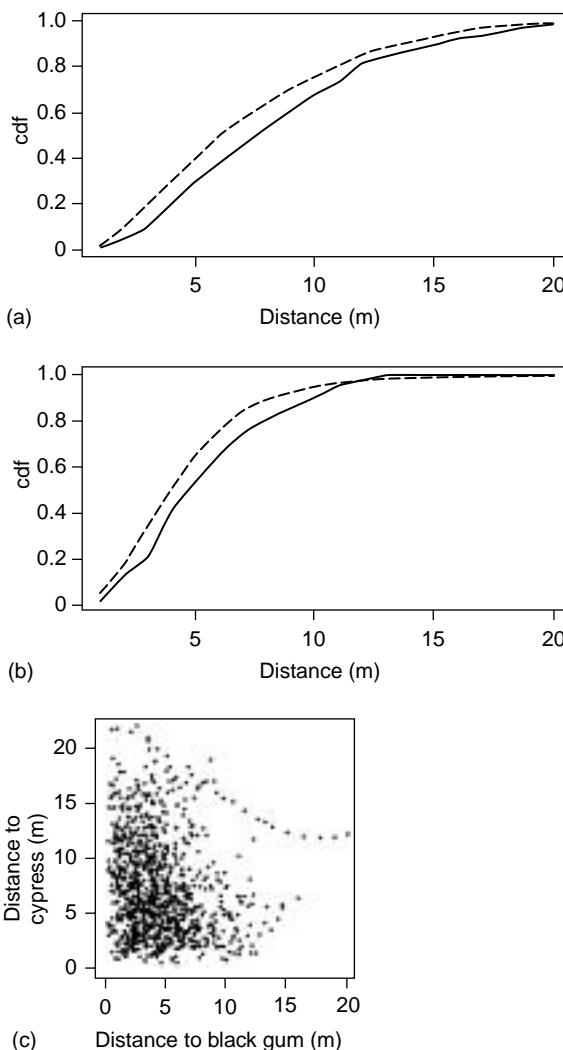

(a)

(b)

(c)

Figure 4 Cdf of point–event distances for all trees and for cypress trees only. (a) Comparison of cdfs of point to cypress tree distances [$F_C(x)$, dotted line] and black gum tree to cypress tree distances [$G_{GC}(x)$, solid line]. (b) Comparison of cdfs of point to black gum tree distances [$F_G(x)$, dotted line] and cypress tree to black gum tree distances [$G_{CG}(x)$, solid line]. (c) Relationship between distances from a randomly chosen point to the nearest cypress tree and nearest black gum tree

pattern. Cypress trees are stochastically further from black gum trees than from randomly chosen points (Figure 4a). Also, black gum trees are stochastically farther from cypress trees than from randomly chosen points (Figure 4b).

The observed differences can be compared with those found under random labeling by using the Kolmogorov–Smirnov two-sample statistics, $\max|F_G(x) - G_{CG}(x)|$ and $\max|F_C(x) - G_{GC}(x)|$, as the test statistics. The observed maximum differences are not unusually large (P value $= 0.109$ for black gum and 0.083 for cypress). The distance from a randomly chosen point to the nearest black gum tree, X_G, is negatively correlated with the distance from the same point to the nearest cypress tree, X_C (Figure 4c; Kendall's $\hat{\tau} = -0.12$, with a one-sided P value, by randomization, of 0.001). This result is consistent with the spatial segregation of the two species. The different results from the three tests are consistent with Diggle and Cox's [27] observation that the correlation test is more powerful than the Kolmogorov–Smirnov test for the sparsely sampled spatial patterns they studied.

The NN contingency table indicates that both species have an excess of NNs of the same species (Table 3). The variances of the cell counts are 38.88 for black gum–black gum and 25.55 for cypress–cypress. The P values can be computed by Monte Carlo randomization or by a normal approximation [30]. In either case, the one-sided P values are small (0.001 or less) for both species.

Nearest Neighbor Methods for Field Experiments

A different set of NN methods can be used to analyze spatially structured field experiments. One of the most common applications is in agronomic variety trials, where many treatments are compared by taking small plots arranged in a rectangular lattice. Traditional methods of controlling for between-plot heterogeneity, such as using a randomized complete block (RCB) design, may not be very effective because the

Table 3 Observed counts (N^O), expected counts (N^E), and z scores for the cypress tree and black gum tree NN contingency table

Species of point	Black gum tree			Cypress tree			Total
	N^O	N^E	z	N^O	N^E	z	
Black gum	149	121.1	4.47	33	60.9	−4.47	182
Cypress	43	60.9	−3.54	48	30.1	3.54	91
Total	192			81			273

large number of treatments forces the blocks to be large. NN methods use information from adjacent plots to adjust for within-block heterogeneity and so provide more precise estimates of treatment means and differences. If there is within-plot heterogeneity on a spatial scale that is larger than a single plot and smaller than the entire block, then yields from adjacent plots will be positively correlated. Information from neighboring plots can be used to reduce or remove the unwanted effect of the spatial heterogeneity and hence improve the estimate of the treatment effect. Data from neighboring plots can also be used to reduce the influence of competition between adjacent plots. Each of these approaches will be briefly discussed below.

Papadakis [49, 50] proposed an analysis of covariance to reduce the effects of small-scale spatial heterogeneity in yields. The value of the covariate for each plot is obtained by averaging **residuals** from the neighboring plots. The choice of neighboring plots depends on the crop, the plot size and shape, and the spacing between plots. In many row crops, the neighboring plots are defined as the two adjacent plots in the row, except that plots at an end of a row have only one neighbor. In other situations, it may be appropriate to consider four neighbors, which include the two between-row neighbors. If the spatial heterogeneity is such that the effects of within-row and between-row neighbors are quite different, then one could compute separate covariates for within-row and between-row neighbors. Once the covariates are computed, treatment effects are re-estimated using an analysis of covariance. For adjustment in one dimension (e.g. along crop rows), the model would be

$$Y_i = \mu + X_i\tau_i + \beta R_i + \varepsilon_i \qquad (15)$$

where Y_i is the observed yield on the ith plot, R_i is the mean residual on neighbors of the ith plot, β is the spatial dependence parameter, μ and τ_i are the parameters in the model for the treatment effects, X_i is the row of the design matrix for the ith plot, and ε_i is the residual variability in yields, which are assumed to be uncorrelated. When β is 0, observations on adjacent plots are independent; larger positive values of β ($\beta < 1$) correspond to increasing spatial correlation between neighboring plots. The observed value of β depends on plot size, plot shape, plot spacing, and the scale of the spatial heterogeneity. Values are often close to 1 when plots are small. In the absence of treatments, and

ignoring edge effects, the Papadakis model implies that correlations between plot yields have a first-order autoregressive structure, with $\mathrm{corr}(Y_i, Y_j) = \lambda^{|i-j|}$, where $\beta = 2\lambda/(1 + \lambda^2)$. Values of β close to 1 imply that λ is also close to 1.

The autoregressive correlation structure implied by the ad hoc Papadakis model is one example of the **random field** approach to a spatially designed experiment [6, 18, 73]. Many other models, including the iterated Papadakis method, the Wilkinson NN [71] model, the Besag and Kempton [9] first-order difference models, the Williams model [69] and the Gleeson–Cullis autoregressive integrated moving average (ARIMA) models [19, 38] (*see* **Time series**) correspond to different specifications of a spatial correlation matrix. Computations are handled either by a general purpose residual maximum likelihood (REML) algorithm for linear **mixed effects** models [e.g. PROC MIXED in **SAS**, lme( ) in **S-PLUS**], or by specialized software for a particular model (e.g. TwoD for two-dimensional Gleeson–Cullis models [37]).

The properties of these methods have been extensively discussed over the past 20 years. Dagnelie [21, 22] provides a relatively recent review and historical summary of the Papadakis model. Wu and Dutilleul [72] use uniformity trial data to compare autoregressive models, difference models, and traditional RCB analyses. The efficiency of a spatial analysis, relative to an RCB design, is usually greater than 1.2 and can be as high as 2 [22]. However, it can give biased estimates of treatment effects [71]. The Papadakis method appears to work best when there are at least three replicates per treatment, many treatments (greater than 10), and strong but patchy spatial heterogeneity [22]. When there is an underlying trend, first-order difference models appear to work well.

Medium-scale spatial heterogeneity usually causes a positive correlation between adjacent plots. When there is competition between plots, neighbors can have a negative effect on the response in adjacent plots [41]. The Papadakis model (15) can be extended to estimate treatment-specific competitive effects. The choice of covariate should be influenced by biological mechanisms. If competition for sunlight is important, a reasonable covariate could be the difference between the mean height of plants in the plot and the mean height on neighboring plots. If disease spread is important, a reasonable covariate could be the mean disease severity on neighboring plots [42]. The coefficient for the covariate [β in (15)] estimates the strength of the competitive relationship.

Experimental design for a study that will use some form of neighbor-adjusted analysis usually focuses on neighbor balance; that is, ensuring that all pairs of treatments occur in adjacent plots equally frequently. Adjacent plots can be defined as only those within the same row (one-dimensional neighbor-balanced designs [70]) or as including both those in the same row and those in the same column (two-dimensional neighbor-balanced designs [34]). The choice will depend on the size, shape, and spacing of the plots and on the biological and physical mechanisms influencing the correlation between plots. Methods for construction of one-dimensional or two-dimensional designs can be found in [33], [34], and [70].

References

[1]   Ashby, E. (1936). Statistical ecology, *Botanical Review* **2**, 221–235.

[2]   Ashby, E. (1948). Statistical ecology. II – a reassessment, *Botanical Review* **14**, 222–234.

[3]   Baddeley, A. & Gill, R.D. (1997). Kaplan–Meier estimators of distance distributions for spatial point processes, *The Annals of Statistics* **25**, 263–292.

[4]   Baddeley, A.J. & Silverman, B.W. (1984). A cautionary example on the use of second-order methods for analyzing point patterns, *Biometrics* **40**, 1089–1093.

[5]   Barnard, G.A. (1963). Contribution to the discussion of Prof. Bartlett's paper, *Journal of the Royal Statistical Society, Series B* **25**, 294.

[6]   Bartlett, M.S. (1978). Nearest neighbour models in the analysis of field experiments (with discussion), *Journal of the Royal Statistical Society, Series B* **40**, 147–174.

[7]   Berman, M. (1986). Testing for spatial association between a point process and another stochastic process, *Applied Statistics* **35**, 54–62.

[8]   Besag, J.E. & Gleaves, J.T. (1973). On the detection of spatial pattern in plant communities, *Bulletin of the International Statistical Institute* **45**, 153–158.

[9]   Besag, J. & Kempton, R. (1986). Statistical analysis of field experiments using neighbouring plots, *Biometrics* **42**, 231–251.

[10] Brewer, R. & McCann, M.T. (1985). Spacing in acorn woodpeckers, *Ecology* **66**, 307–308.

[11] Brown (1975). A test of randomness in nest spacing, *Wildlife* **26**, 102–103.

[12] Brown, D. & Rothery, P. (1978). Randomness and local regularity of points in a plane, *Biometrika* **65**, 115–122.

[13] Byth, K. (1982). On robust distance-based intensity estimators, *Biometrics* **38**, 127–135.

[14] Byth, K. & Ripley, B.D. (1980). On sampling spatial patterns by distance methods, *Biometrics* **36**, 279–284.

[15] Clark, P.J. & Evans, F.C. (1954). Distance to nearest neighbor as a measure of spatial relationships in populations, *Ecology* **35**, 445–453.

[16] Cox, T.F. (1981). Reflexive nearest neighbours, *Biometrics* **37**, 367–369.

[17] Cox, T.F. & Lewis, T. (1976). A conditioned distance ratio method for analyzing spatial patterns, *Biometrika* **63**, 483–491.

[18] Cressie, N. (1991). *Statistics for Spatial Data*, Wiley, New York.

[19] Cullis, B.R. & Gleeson, A.C. (1991). Spatial analysis of field experiments – extension to two dimensions, *Biometrics* **47**, 1449–1460.

[20] Cuzick, J. & Edwards, R. (1990). Spatial clustering for inhomogeneous populations (with discussion), *Journal of the Royal Statistical Society, Series B* **52**, 73–104.

[21] Dagnelie, P. (1987). La méthode de Papadakis en expérimentation agronomique: considérations historiques et bibliographiques, *Biométrie et Praximétrie* **27**, 49–64.

[22] Dagnelie, P. (1989). The method of Papadakis in agricultural estimation: an overview, *Biuletyn Oceny Odmian* **21**, 111–122.

[23] Devroye, L. & Györfi, L. (1985). *Nonparametric Density Estimation. The L_1 View*, Wiley, New York.

[24] Diggle, P.J. (1979). On parameter estimation and goodness-of-fit testing for spatial point patterns, *Biometrics* **35**, 87–101.

[25] Diggle, P.J. (1983). *Statistical Analysis of Spatial Point Patterns*, Academic Press, London.

[26] Diggle, P.J. & Chetwynd, A.G. (1991). Second-order analysis of spatial clustering for inhomogeneous populations, *Biometrics* **47**, 1155–1163.

[27] Diggle, P.J. & Cox, T.F. (1983). Some distance-based tests of independence for sparsely-sampled multivariate spatial point patterns, *International Statistical Review* **51**, 11–23.

[28] Diggle, P.J., Besag, J. & Gleaves, J.T. (1976). Statistical analysis of spatial point patterns by means of distance methods, *Biometrics* **32**, 659–667.

[29] Dixon, P.M. (1992). Testing spatial independence in multivariate point processes, Institute of Statistics Mimeo Series #2215, Raleigh.

[30] Dixon, P.M. (1994). Testing spatial segregation using a nearest-neighbor contingency table, *Ecology* **75**, 1940–1948.

[31] Donnelly, K.P. (1978). Simulations to determine the variance and edge effect of total nearest-neighbour distance, in *Simulation Studies in Archaeology*, I. Hodder, ed., Cambridge University Press, Cambridge, pp. 91–95.

[32] Everitt, B.S. (1993). *Cluster Analysis*, 3rd Edition, Edward Arnold, London.

[33] Federer, W.T. & Basford, K.E. (1991). Competing effects designs and models for two-dimensional field arrangements, *Biometrics* **47**, 1461–1472.

[34] Freeman, G.H. (1979). Some two-dimensional designs balanced for nearest neighbours, *Journal of the Royal Statistical Society, Series B* **41**, 88–95.

[35] Friedman, J., Bentley, J.L. & Finkel, R.A. (1977). An algorithm for finding best matches in logarithmic expected time, *ACM Transactions on Mathematical Software* **3**, 209–226.

[36] Gignoux, J., Duby, C. & Barot, S. (1999). Comparing the performance of Diggle's tests of spatial randomness for small samples with and without edge-effect correction: application to ecological data, *Biometrics* **55**, 156–164.

[37] Gilmour, A.R. (1992). TwoD: a program to fit a mixed linear model with two dimensional spatial adjustment for local trend, NSW Agriculture Research Center, Tamworth, Australia.

[38] Gleeson, A.C. & Cullis, B.R. (1987). Residual maximum likelihood (REML) estimation of a neighbour model for field experiments, *Biometrics* **43**, 277–288.

[39] Goodall, D.W. (1965). Plot-less tests of interspecific association, *Journal of Ecology* **53**, 197–210.

[40] Hodder, I. & Orton, C. (1976). *Spatial Analysis in Archaeology*, Cambridge University Press, Cambridge.

[41] Kempton, R. (1982). Adjustment for competition between varieties in plant breeding trials, *Journal of Agricultural Science* **98**, 599–611.

[42] Kempton, R. (1991). Interference in agricultural experiments, *Proceedings of the Second Biometric Society East/Central/Southern African Network Meeting*, Harare, Zimbabwe.

[43] Lawson, A. (1988). On tests for spatial trend in a non-homogeneous Poisson process, *Journal of Applied Statistics* **15**, 225–234.

[44] Lotwick, H.W. & Silverman, B.W. (1982). Methods for analyzing spatial processes of several types of points, *Journal of the Royal Statistical Society, Series B* **44**, 406–413.

[45] Mathsoft (1996). *S + SPATIALSTATS User's Manual, Version 1.0*, Mathsoft Inc., Seattle.

[46] Moore, P.G. (1954). Spacing in plant populations, *Ecology* **35**, 222–227.

[47] Murtagh, F. (1984). A review of fast techniques for nearest neighbor searching, *Compstat* **1984**, 143–147.

[48] Neyman, J. & Scott, E.L. (1952). A theory of the spatial distribution of galaxies, *Astrophysical Journal* **116**, 144–163.

[49] Papadakis, J.S. (1973). Méthode statistique pour des expériences sur champ, Institut d'Amélioration des Plantes a Thessaloniki (Thessalonika, Greece), *Bulletin Scientifique*, No. 23.

[50] Papadakis, J.S. (1984). Advances in the analysis of field experiments, *Proceedings of the Academy of Athens* **59**, 326–342.

[51] Pickard, D.K. (1982). Isolated nearest neighbors, *Journal of Applied Probability* **19**, 444–449.

[52] Pielou, E.C. (1961). Segregation and symmetry in two-species populations as studied by nearest-neighbour relationships, *Journal of Ecology* **49**, 255–269.

[53] Pielou, E.C. (1977). *Mathematical Ecology*, Wiley, New York.

[54] Pollard, J.H. (1971). On distance estimators of density in randomly distributed forestes, *Biometrics* **27**, 991–1002.

[55] Rathbun, S.L. (1996). Estimation of Poisson intensity using partially observed concomitant variables, *Biometrics* **52**, 226–242.

[56] Ripley, B.D. (1977). Modelling spatial patterns (with discussion), *Journal of the Royal Statistical Society, Series B* **39**, 172–212.

[57] Ripley, B.D. (1979). Tests of 'randomness' for spatial point patterns, *Journal of the Royal Statistical Society, Series B* **41**, 368–374.

[58] Ripley, B.D. (1981). *Spatial Statistics*, Wiley, New York.

[59] Ripley, B.D. (1996). *Pattern Recognition and Neural Networks*, Cambridge University Press, Cambridge.

[60] Rowlinson, B.S. & Diggle, P.J. (1992). Splancs: Spatial Point Pattern Analysis Code in S-Plus, Technical Report 92/63, Lancaster University, Lancaster, UK.

[61] Schilling, M.F. (1986). Mutual and shared neighbor probabilities: finite- and infinite-dimensional results, *Advances in Applied Probability* **18**, 388–405.

[62] Silverman, B. & Brown, T. (1978). Short distance, flat triangles, and Poisson limits, *Journal of Applied Probability* **15**, 815–825.

[63] Simberloff, D. (1979). Nearest neighbor assessments of spatial configurations of circles rather than points, *Ecology* **60**, 679–685.

[64] Stoyan, D. & Penttinen, A. (2000). Recent applications of point process methods in forestry statistics, *Statistical Science* **15**, 61–78.

[65] Stoyan, D. & Stoyan, H. (1994). *Fractals, Random Shapes and Point Fields: Methods of Geometrical Statistics*, Wiley, Chichester.

[66] Upton, G.J.G. & Fingleton, B. (1985). *Spatial Data Analysis by Example*, Vol. 1: *Point Pattern and Quantitative Data*, Wiley, Chichester.

[67] Venables, W.N. & Ripley, B.D. (1994). *Modern Applied Statistics with S-Plus*, Springer-Verlag, New York.

[68] Waller, L.A., Turnbull, B.W., Clark, L.C. & Nasca, P. (1994). Spatial pattern analyses to detect rare disease clusters, in *Case Studies in Biometry*, N. Lange et al. eds, Wiley, New York, pp. 3–23.

[69] Williams, E.R. (1986). A neighbour model for field experiments, *Biometrika* **73**, 279–287.

[70] Williams, R.M. (1952). Experimental designs for serially correlated observations, *Biometrika* **39**, 151–167.

[71] Wilkinson, G.N., Eckert, S.R., Hancock, T.W. & Mayo, O. (1983). Nearest neighbour (NN) analysis of field experiments (with discussion), *Journal of the Royal Statistical Society, Series B* **45**, 151–211.

[72] Wu, T. & Dutilleul, P. (1999). Validity and efficiency of neighbor analyses in comparison with classical complete and incomplete block analyses of field experiments, *Agronomy Journal* **91**, 721–731.

[73] Zimmerman, D.L. & Harville, D.A. (1991). A random field approach to the analysis of field-plot experiments and other spatial experiments, *Biometrics* **47**, 223–239.

(*See also* **Point processes, spatial–temporal**; **Spatial analysis in ecology**; **Spatial design, optimal**)

PHILIP M. DIXON