*Ecology*, 79(6), 1998, pp. 2180–2192


# A NEW METHOD FOR DETECTING SPECIES ASSOCIATIONS WITH SPATIALLY AUTOCORRELATED DATA

STEPHEN H. ROXBURGH AND PETER CHESSON

*Ecosystem Dynamics Group, Research School of Biological Sciences, Institute of Advanced Studies, Australian National University, Canberra, Australian Capital Territories 0200, Australia*

*Abstract.* Many organisms display patchiness in their distribution patterns over a wide range of spatial scales. Patchy distribution patterns can be caused by processes such as growth, migration, reproduction, and mortality, which result in neighboring areas being more likely to contain a species than distant areas, a phenomenon known as positive spatial autocorrelation. When species are patchily distributed, the within-species spatial randomness assumptions of the standard statistical tests for detecting species associations are seriously violated. Using these tests under such circumstances can lead to incorrect rejection of the null hypothesis. To address this problem we introduce a new test for detecting species associations—the random patterns test. This test takes into account spatial autocorrelation by including the characteristics of the spatial pattern of each species into the null model. A randomization procedure was used to generate the null distribution of the test statistic. The random patterns test is illustrated with data collected from an herbaceous understory community of a *Eucalyptus* forest near Canberra, Australia.

*Key words: pattern analysis; randomization test; semivariograms; spatial autocorrelation; species association tests.*

## INTRODUCTION

Interspecific associations arise when two or more species co-occur either more or less frequently than expected due to chance alone. Positive associations between two species can occur when both select the same habitat or have the same environmental requirements. Conversely, negative associations can occur if the species have differing ecological requirements (Dale 1977). Association, in either the positive or negative direction, can also occur as a direct consequence of biotic interactions such as mutualism, competition and predation. Although it is not possible to unambiguously infer the action of specific processes from the examination of patterns alone (Schluter 1984, Rejmánek and Lepš 1996), association analyses remain a valuable tool for ecologists. Such analyses can be used for generating hypotheses about the factors responsible for the patterns, and hence can be used for identifying particular patterns that may be worthy of further study.

A number of statistical tests have been utilized for detecting species associations. These include correlation analysis (Greig-Smith 1983, O'Connor and Aarssen 1987, Myster and Pickett 1992), analysis by contingency table (Greig-Smith 1983, Dale et al. 1991), variance tests (Schluter 1984, McCulloch 1985), and the use of cross-variograms (Rossi et al. 1992). One major limitation of these tests is that they assume, for each species, within-species randomness of the spatial distribution patterns. Throughout this paper we use the term ''within-species randomness'' to refer to the situation where the occurrence of a species in one quadrat (or other sampling unit) is independent of the occurrence of the same species in other quadrats, regardless of the relative spatial locations of the quadrats.

Manuscript received 4 November 1996; revised 24 June 1997; accepted 15 July 1997; final version received 7 August 1997.

Because species are typically distributed nonrandomly in space, the within-species randomness assumption is often violated. For example, two of the most common nonrandom distribution patterns are the aggregation of a species into patches, and gradients. Patchy distributions can result from processes such as growth, migration, reproduction, mortality and natural selection, and can also arise when suitable habitats are distributed in a patchy manner. Gradients in abundance are often associated with gradual environmental changes, such as the change in environmental conditions with increasing altitude (e.g., Whittaker and Niering 1975).

Patchy or clumped patterns, where an observation in one area makes it more likely that the same observation will also be made in neighboring areas, are said to exhibit positive spatial autocorrelation. With traditional tests of interspecific association the presence of positive spatial autocorrelation violates the assumption of within-species randomness, and results in an elevated Type I error rate, i.e., an increase in the risk of concluding a test statistic significant, even when the species are unassociated (Tavaré and Altham 1983, Dale et al. 1991, Legendre 1993, Palmer and van der Maarel 1995). This feature of these tests is illustrated in the next section.

Because the distribution patterns of most organisms

are inherently nonrandom, then the problems caused by spatial autocorrelation are very general indeed. More specifically, traditional association analyses are valid only if it can be confirmed that, for all species, the within-species patterns are random. If this cannot be demonstrated then any results derived from such analyses must be treated with a great deal of caution.

The purpose of this paper is to describe a new statistical test for detecting species associations which is valid even when the assumption of within-species spatial randomness is violated. Our method uses a Monte Carlo technique that generates null distribution patterns that are spatially autocorrelated within species, but independent between species. The observed patterns can then be compared with these null patterns for a true assessment of the statistical significance of species association, free of the assumption of within-species randomness.

## The Problem of Spatial Autocorrelation

Fig. 1a shows the spatial patterns of two species, *Poa sieberiana* and *Lagenifera stipitata,* within a 40-cm square experimental plot taken from our larger field experiment (see *Application of the random patterns test*: *Site description and sampling methods* for further details of the experiment). Both species show significant positive spatial autocorrelation as measured by Moran's *I,* with neighbors defined as for the Queen's case (Goodchild 1986) (*Poa*: Moran's $I = 0.694$, $P < 0.001$; *Lagenifera*: Moran's $I = 0.290$, $P < 0.001$). At this sampling scale, the patchiness exhibited by these species is caused by a combination of the clonal growth habit of these plants, and the fact that our smallest sampling unit is less than the average size of an individual leaf (for *Poa*) or ramet (for *Lagenifera*).

To demonstrate how spatial autocorrelation invalidates traditional statistical techniques, the null hypothesis that *Poa* and *Lagenifera* occur within this plot independently of one another was tested through the use of the standard $2 \times 2$ contingency table analysis (Ludwig and Reynolds 1988). In Fig. 1a the species were observed to co-occur in 77 out of 784 cells. The expected number of cells to contain both species, calculated under the assumption of independence between species, is 63.25 (Table 1). Based on the standard contingency table chi-squared test, which assumes randomness within species as well as independence between species, this overlap is significantly ($P = 0.003$) less than the observed value of 77 and the conclusion is that *Poa* and *Lagenifera* are positively associated.

However, as the original data show spatial autocorrelation for both species, the results of this test are invalid. This is illustrated in Fig. 1b. Each of these three plots contain the same number of cells of *Poa* (506) and the same number of cells of *Lagenifera* (98) as in the observed, however each species has been assigned to the 784 cells under the assumption of within-species randomness. This assumption results in a problem related to pseudo-replication (Hurlbert 1984). In the test the individual cells are taken as independent sampling units, whereas within-species patchiness makes it clear that the cells are not independent. Statistically we can say there are fewer degrees of freedom when the pattern is autocorrelated, because each new record in a cell does not correspond to a new and independent piece of information. The result is a *P* value that is too low leading to an inflated estimate of significance of the association. It is also clear from the patterns in Fig. 1b that the independent assignment of species to cells destroys the original spatial pattern of each species, with the result that the randomized patterns in Fig. 1b are biologically unrealistic, i.e., in nature the spatial patterns of *Poa* and *Lagenifera,* at the 40-cm scale, are never observed to be as fragmented as those depicted. A biologically sensible null model would therefore take the spatial patterns of each species into account before testing for a significant association.

## Association Analyses by Randomization Test

Before describing our solution to the spatial autocorrelation problem, and to introduce the concept of repeated sampling from a null model, the above statistical test can be repeated using the random maps in Fig. 1b as the basis of a randomization test (Fig. 2).

To perform any sort of statistical test three steps are required. Firstly, the null hypothesis should be stated. Secondly, a test statistic should be chosen which is able to discriminate between the null and alternative hypotheses. Thirdly, a null model must be constructed which is able to generate a probability distribution of the test statistic under the assumption that the null hypothesis is true (Noreen 1989).

For the test of interspecific association the null hypothesis is that *Poa* and *Lagenifera* are distributed independently of each other. A possible test statistic is the chi-squared, as defined in Table 1. However this statistic is simply a function of the observed overlap, and an equivalent but simpler method is to regard the observed overlap, i.e., the number of cells containing both species, as the test statistic. Recall from the example pattern that the observed overlap was 77 cells (Fig. 1a).

For the randomization test the *P* value is estimated from repeated sampling from the null model. In this case the null model is straightforward. Each species is allocated independently and at random among the 784 ($28 \times 28$) cells with a frequency equal to that in the observed plot, as in Fig. 1b. The overlap (the test statistic) is calculated for each random rearrangement, and this process is repeated for a large number of such random permutations of the data to construct the null probability distribution of the test statistic. Fig. 1b shows three possible permutations. The *P* value is estimated by comparing the observed value of the test statistic (77) with the probability distribution of the test statistic derived from 1000 random permutations,

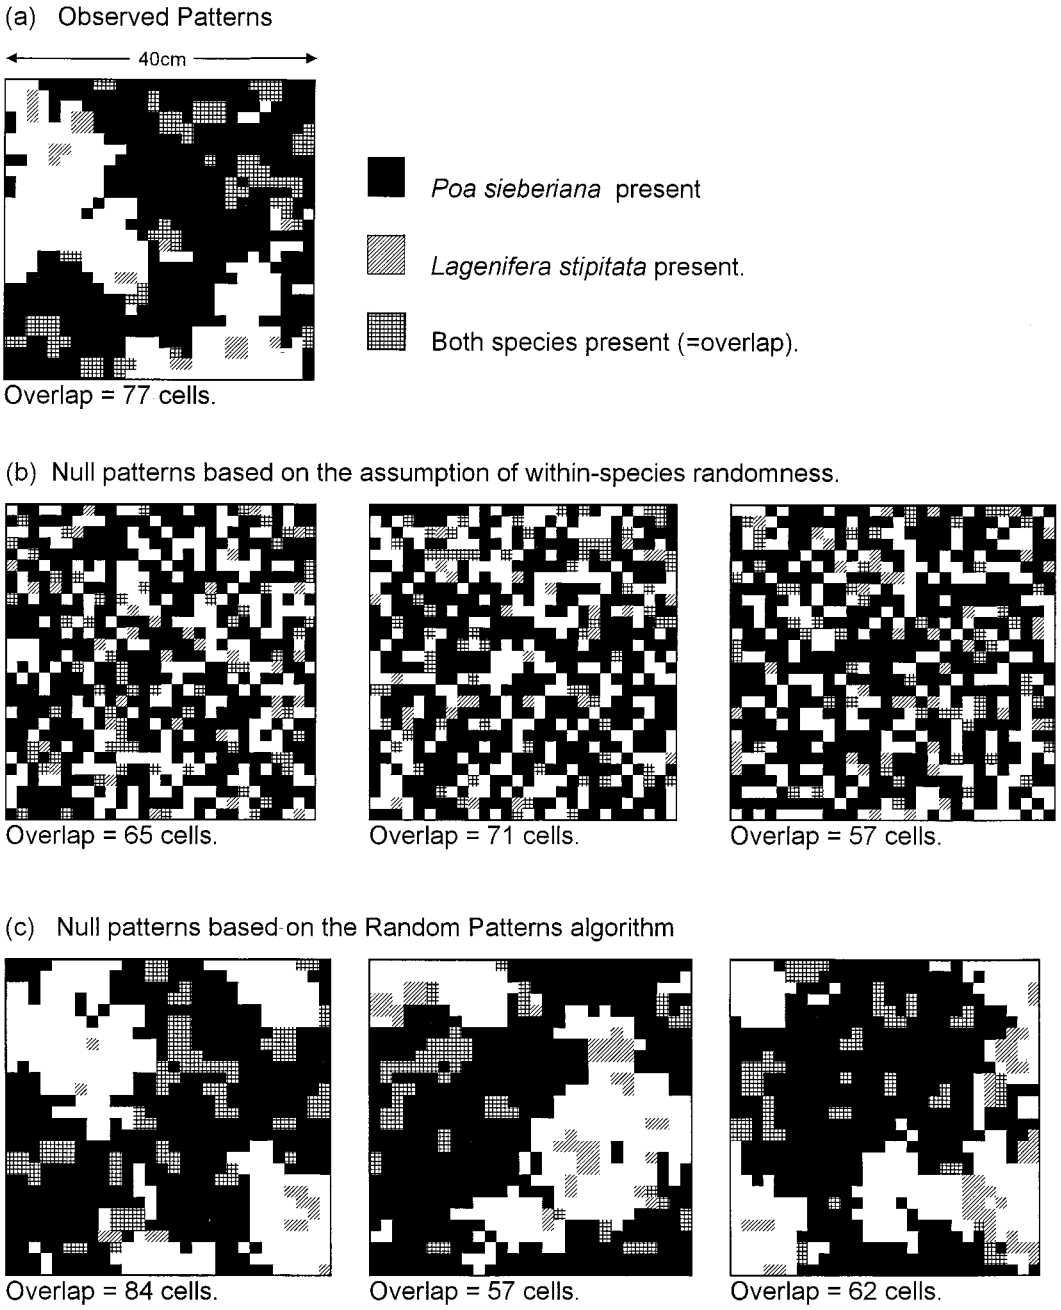


FIG. 1. An example of the overlap between two species exhibiting positive spatial autocorrelation. (a) The observed spatial patterns of *Poa sieberiana* and *Lagenifera stipitata* within a 40 cm square plot subdivided into a grid of 28 × 28 cells. (b) Three random (null) patterns generated under the chi-squared assumption of within-species spatial randomness. Each species is assigned to the cells at an overall frequency equal to that in the observed plot, but independent of the presence of the same species in neighboring cells. (c) Three random (null) patterns generated using the random patterns method.

and was found to be 0.002 (Table 2), which is very close to the *P* value estimate taken from chi-squared tables. Note that the estimates of the *P* values derived from the randomization test and the chi-squared tables are not exactly the same. This is because the randomization test was based on only 1000 randomizations, which is an approximation to the analytical result (0.003) you would get if you could perform an infinite number of randomizations.

## Previous Methods to Address the Effects of Spatial Autocorrelation

Before describing the random patterns test, a brief survey of methods that have been previously proposed

TABLE 1. Chi-squared test of association between *Lagenifera stipitata* and *Poa sieberiana* from the data in Fig. 1a.

| | *Lagenifera* | | | |
|---|---|---|---|---|
| | Present | | Absent | |
| *Poa* | Obs. | Exp. | Obs. | Exp. |
| Present | 77 | 63.25 | 429 | 442.75 |
| Absent | 21 | 34.75 | 257 | 243.25 |

*Note:* $\chi^2 = 9.634$, for 1 df, $P(x > \chi^2) = 0.003$

to address the problems of spatial autocorrelation will be given.

### *Adjustment methods*

These methods do not alter the calculation of the test statistic or the mechanics of the test, but rather the test statistic is adjusted a posteriori, resulting in a new estimate which reflects the correct Type I error rate. For the two-way contingency table analysis Tavaré and Altham (1983) quantified the effect that various degrees of serially dependent data, generated by Markov dependent sequences, has on increasing the Type I error rate. From this they derived deflation factors which were then applied to the calculated test statistic. By deflating the magnitude of the test statistic by the appropriate amount, depending on the degree of nonrandomness in the data, an estimate of the significance of the test statistic in the presence of autocorrelation is obtained. A limitation of Tavaré and Altham's (1983) method is that it is valid only when the within-species spatial patterns meet Markovian assumptions (Dale et al. 1991).

Dale et al. (1991) extended Tavaré and Altham's (1983) approach and suggested a Monte Carlo procedure that can be applied to the analysis of multispecies $2^k$ contingency tables, where $k$ is the number of species. Their method involves defining a null model which includes information on the spatial patterns of each species. This null model is used to generate a null distribution of the test statistic, in their example the $G$ statistic. This distribution is then used to calculate the appropriate deflation factor to give the appropriate Type I error rate. This deflation factor is then applied to the observed $G$ statistic, which is then compared with the chi-squared distribution to calculate the final $P$ value.

However, the calculation of the deflation factor is an unnecessary complication, as the significance of the observed test statistic can be obtained by direct comparison with the distribution of $G$ values derived from the null model, i.e., by performing a randomization test. If the observed test statistic is unusual, i.e., significant, then it will be found in one of the extreme tails of the distribution of values generated under the null model. By determining the significance in this way allows a wider choice of test statistics to be defined, thus eliminating the need to approximate to a known distribution such as the chi-square.

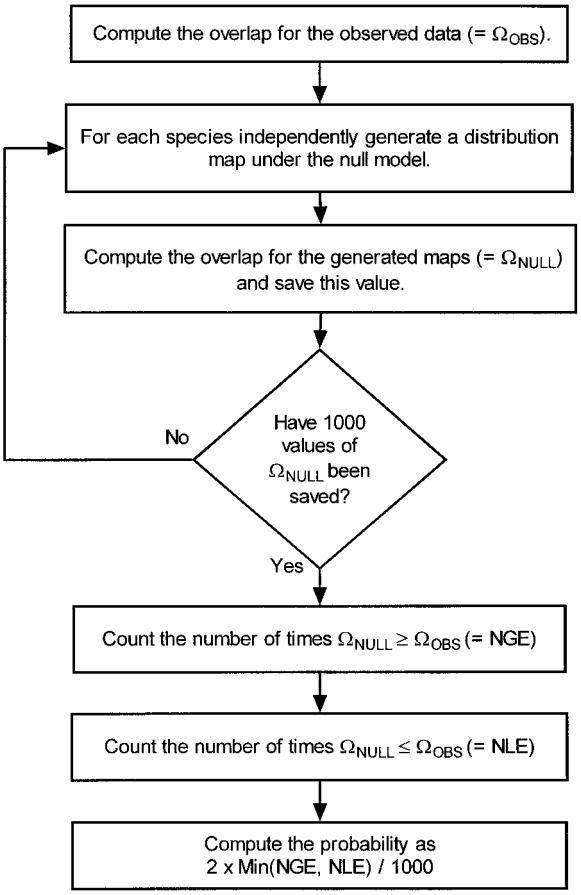


FIG. 2. Generalized flow chart for a randomization test used to detect pairwise interspecific association, based on $n = 1000$ randomizations. NGE (NLE) is the number of test statistics calculated from the random maps that were greater than (less than) or equal to the test statistic calculated from the observed data. The $P$ value is calculated as the minimum of (NGE, NLE) divided by $n$, then multiplied by 2 to effect a two-tailed test.

### *Algorithms for simulating spatial structure*

The chi-squared test by randomization, described in the previous section, generates null patterns for each species independently, but with the added assumption of randomness within species. It is this within-species randomness assumption that results in fragmentation of the observed spatial patterns, leading to a biased Type I error rate. To overcome these difficulties, several algorithms exist that simulate spatial patterns by incorporating information on the observed within-species spatial structure.

McElhany et al. (1995) suggested an algorithm called the additive method, which is similar to the random patterns approach we describe below. The additive method builds up a pattern cell by cell, such that the locations of new filled cells produce a pattern that ex-

TABLE 2. Randomization test results. The expected values are the average of the 1000 overlap values generated under each null model. The ''Random assignment'' column shows the results for a randomization test equivalent to the chi-squared test shown in Table 1. The ''Random patterns test'' column shows the results for a test using the random patterns method applied to the same data.

| | Random assignment | Random patterns test |
|---|---|---|
| Observed overlap | 77 | 77 |
| Expected overlap | 62.7 | 63.2 |
| Number of random maps with an overlap $\geq$ observed | 1 | 41 |
| *P* | 0.002 | 0.082 |

hibits the same degree of clumpiness as in the observed data. Although capable of reproducing random spatial patterns very similar to those in the observed data, this method does have the disadvantage of being computationally inefficient when reconstructing large spatial arrays (Real and McElhany 1996).

Watkins and Wilson's (1992) patch model randomly allocates species to cells, similar to that under the independence assumption in Fig. 1b, but only locally, i.e., the decision on whether to assign a species to a cell is based on the presence of that species in cells that occur within a set distance from the target cell, defined by the ''patch'' size. In the context of a two-species association test, such local reallocation of the species would control for within-plot variation in density, for example where high density patches within a plot occur due to local dispersal, or clustering around pockets of high resource concentration (the ''waterhole'' effect, Pielou 1977).

It is important to recognize that the assumptions behind the patch model are quite different from those of the other methods discussed in this section. When incorporated into a statistical test, all of the other methods are designed to detect any departure from randomness, regardless of the underlying mechanism responsible for the departure. Deviation from randomness can be caused by either the species growing in different environments within the sampled area (beta-niche differences), or as a result of biotic interactions such as competition and mutualism (alpha-niche differences). Once a nonrandom pattern is detected, then it is up to the researcher to determine what possible mechanisms could be responsible. In contrast, the patch model is an attempt to reduce the chances of detecting associations that are due to beta-niche differentiation, leaving as a residual those associations which are the result of interspecific interactions (Wilson 1995).

Palmer and van der Maarel's (1995) rotation/reflection and random shifts methods construct spatial patterns at random through transformations of the observed spatial array. With the rotation/reflection method, the overall spatial pattern of each species is retained as a fixed aspect of the null distribution. Species patterns in space are generated by having each species independently undergo a rotation and reflection of the whole plot. The rotated and/or reflected patterns are then reassembled and the test statistic calculated. Although this method does retain the observed spatial characteristics of each species, an unfortunate result is that the central quadrats for each species do not move, and therefore are not adequately randomized, resulting in a loss of statistical power (Manly 1991).

The random shifts method combines a rotation/reflection with a further transformation. With a raster such as Fig. 1a this transformation is a random shift in both the *x*- and *y*-axes. When the pattern is shifted beyond the edge, it is automatically wrapped around to the opposite side. The main disadvantage of this method is the disruption to the overall spatial characteristics of the pattern, the extent of the disruption depending on the particular characteristics of the observed spatial pattern (Palmer and van der Maarel 1995). This occurs particularly when sharp boundaries at the edge of the sampled area appear in the center of the plot following a random shift, creating an unfortunate artifact. This aspect of the random shifts algorithm is investigated in greater detail below. Note that the rotation/reflection and random shifts tests, contrary to Palmer and van der Maarel's claim, are not alternatives to the patch model, as there is no attempt in these tests to separate associations due to environmental heterogeneity from those due to interspecific interactions (Wilson 1995).

Several other techniques also exist for generating spatial patterns, although space does not permit a full discussion of them all. For further information the reader is referred to the geostatistical literature (e.g., Wackernagel 1995, Cressie 1993, Journal and Huijbregts 1978); the review by Real and McElhany (1996) covering a range of methods, including the use of matrix methods for generating spatial patterns; Mangel and Adler's (1994) ''force to be full'' method based on the concept of ''structure functions''; and the method employing a fractal landscape generating algorithm described by Palmer (1992).

### THE RANDOM PATTERNS TEST

The random patterns test was developed to overcome the difficulties associated with some of the above methods. In particular, the rotation/reflection and random shifts methods both aim to generate null patterns in which the within-species clumping is retained while introducing statistical independence between species. However, with the rotation/reflection method applied to a test of association between two species there are too few null patterns for a test with adequate statistical power. With the random shifts method there are enough null patterns to make a test, but unrealistic patterns can occur.

The random patterns test uses a different method of null pattern generation to overcome these difficulties.

The basic assumption is that it should be possible to adequately characterize the within-species patchiness numerically, using ''clumping'' statistics. Any two within-species patterns with the same values for these statistics are presumed to be equally good representations of the outcome of whatever process generated the within-species patterns. Thus, the random patterns test has two important parts. First, derivation of clumping statistics capable of characterizing the within-species patterns. Second, a method of randomly generating within-species patterns with the same values of the clumping statistics as the observed pattern. By independent generation of such clumped patterns for the different species, a null distribution of species overlap can be obtained.

*Statistical characterization of the observed patterns*

Initial attempts to reconstruct the spatial patterns for each species using classical spatial autocorrelation statistics such as Moran's *I* and Geary's *c* resulted in random patterns which often did not resemble those observed. For this reason we derived four alternative statistics on which to base the analysis, each of which quantifies a particular aspect of the patterns. These were the number of edge contacts (*E*), corner contacts (*C*), open areas (*O*), and solid areas (*S*). These statistics are defined in the Appendix.

*Generation of random patterns from the observed statistics*

To randomly generate patterns with the spatial characteristics of the observed patterns we firstly assigned the species to the cells under the assumption of within-species randomness (Fig. 1b). To test agreement with the observed within-species clumping we then calculated

$$\phi = \left|\frac{E_{\text{NULL}}}{E_{\text{OBS}}} - 1\right| + \left|\frac{C_{\text{NULL}}}{C_{\text{OBS}}} - 1\right| + \left|\frac{O_{\text{NULL}}}{O_{\text{OBS}}} - 1\right| + \left|\frac{S_{\text{NULL}}}{S_{\text{OBS}}} - 1\right|$$

where $E_{\text{OBS}}$ is the edge contacts (*E*) for the observed pattern, and $E_{\text{NULL}}$ the edge contacts for the initial random arrangement of cells, and so on. From this equation it can be seen that when the spatial pattern for the randomized map is the same as that in the observed, then the function $\phi$ is at its minimum of zero. To replicate the observed spatial characteristics from the initial randomization therefore required minimizing the function $\phi$. This was achieved by swapping two cells at random (one occupied, and one not occupied) and recalculating $\phi$. If $\phi$ decreased, i.e., the spatial pattern in the randomized pattern became more similar to the observed, then the swap was retained, otherwise the two cells were returned to their original state, and another pair selected. This process was repeated until $\phi$ reached a predetermined limit, $\alpha$, which is set close to zero to ensure a close match between the random and observed patterns. For the randomizations presented in this paper $\alpha = 0.01$, which was found to produce random patterns very similar to those observed, while at the same time significantly reducing the computing time required (results not presented).

For any observed pattern with a reasonable number of filled and empty cells there are an enormous number of such randomly generated clumped patterns, thus dealing with the power limitations of the rotation/reflection method. Additionally, the randomly generated patterns all share the same within-species statistical properties, addressing the problems associated with the random shifts method.

*Implementation of the random patterns test*

The implementation of the random patterns test is very similar to the randomization approach to the chi-squared test, as summarized in Fig. 2. Firstly, a null species distribution pattern is generated independently for each species using the procedure described immediately above. These patterns are then combined and the between-species overlap (the test statistic) calculated. This randomization procedure is repeated *n* times to create the null distribution of the test statistic. The observed overlap value is then compared with this distribution, and the *P* value calculated as a two-tailed test (Fig. 2). When conducting a randomization test, the more randomizations that are performed, the greater the accuracy of the estimated *P* value. We chose $n = 1000$ randomizations for all tests in this paper, as this provides an acceptable level of statistical power (Edgington 1995), whilst at the same time restricting computing time to within manageable limits. Indeed, implementing the random patterns procedure is extremely computer intensive. In the Appendix are described the methods used in the pattern generating algorithm to increase computer efficiency.

Reanalysis of the test for association between *Poa* and *Lagenifera* in Fig. 1a using the random patterns test highlights the liberal nature of the standard contingency table analysis. In contrast to the conclusions from the contingency table analysis, which detected a strong positive relationship between the two species ($P = 0.003$, Table 1), the random patterns test suggests, however, that these data contain no strong evidence of an association between the two species ($P = 0.082$, Table 2). Note that incorporating the spatial characteristics of each species into the random patterns null model produces patterns that are now 'typical' of those observed in the actual community, which is in contrast to the biologically unreasonable patterns produced under the assumption of within-species randomness (compare Figs. 1b and 1c with 1a).

## Confirmation of the Random Patterns Test

Before applying the random patterns test to real data it was first necessary to confirm that (a) the random

patterns algorithm did indeed recreate the spatial characteristics of the observed patterns, and (b) the random patterns test was generating the correct type I error rate.

### *Semivariogram analysis of randomly generated patterns*

Patterns generated by the random patterns algorithm visually agreed well with the observed single species patterns, but rather than rely on this subjective technique we checked the patterns against another set of statistics not used in the specification of the null model. To achieve this we compared the characteristics of the observed pattern with those generated under the random patterns method, by comparing semivariograms between the observed and randomized patterns (Burrough 1987). Ten species distribution patterns were analyzed, chosen to include a range of the spatial patterns observed in our data. For each of the 10 patterns a semivariogram was constructed for the observed data, and this was compared to semivariograms derived from 50 patterns generated by the random patterns method, and 50 patterns generated by Palmer and van der Maarel's (1995) random shifts method. The random shifts method was included to investigate the degree to which this test disrupts the overall spatial characteristics the patterns. Note that the statistics used to quantify the spatial patterns ($E$, $C$, $O$, and $S$) characterize only the first order (distance = 1) nature of the patterns, hence we would expect the semivariance of the variograms at distance =1 to be very similar between the observed and randomized patterns. However the semivariogram approach additionally describes the higher order components of the pattern (distances >1) not included in the random patterns procedure, and hence provides a valid method for comparison.

Semivariograms for a variety of spatial patterns for *Poa* and *Lagenifera*, for both the random patterns and random shifts tests, are displayed in Figs. 3 and 4. Although the patterns tested were selected from only two species, the distributions of *Lagenifera* are typical of the range of spatial patterns characteristic of the remaining species in this community. Note that a semivariogram for a single clumped pattern necessarily shows high fluctuations when the lag size (distance) becomes large, about 1/2 the plot size in our plots. This does not reflect some genuine property of clumping, but imprecision in estimation of the semivariogram at large distances due to few data at these distances. This is true of our randomly generated patterns too, but not of the average of the semivariograms of many independently generated patterns, as one would expect. To illustrate this feature, Figs. 3a and 4a show individual semivariograms for 50 patterns generated by the random patterns and random shifts algorithms from the observed species maps shown in Figs. 3b and 4b, respectively. However, to enable a direct comparison between the random patterns and random shifts methods, Figs. 3b–f and 4b–f show only the mean of the 50 randomly generated patterns.

For the random patterns test, the semivariograms for the observed *Poa* and *Lagenifera* distributions were predominantly within the bounds of the 50 randomly generated patterns, confirming that the test indeed generates ''typical'' spatial arrangements of the observed patterns (Figs. 3 and 4). We are therefore confident that the random patterns test is capable of producing randomized distributions which accurately characterize the overall spatial patterns of the species.

In contrast, inspection of the random *Poa* distributions generated using the random shifts method showed evidence for systematic bias. All five random shifts semivariograms show consistent over-estimation of the semivariance at the smaller distances, as indicated by the open circles in Fig. 3b–f. This shows that, for areas in close proximity, the spatial patterns generated by the random shifts test are consistently less spatially autocorrelated than in the observed patterns. Also, there was consistent under-estimation of the semivariance at the larger distances, dramatically so in the case of Fig. 3e. These features are a result of the fragmentation of the large patches of *Poa* due to the ''wrap around'' effect of the random shifts procedure, resulting in distributions which are more random than reality when viewed at the small distances, and less random at the larger. As already noted, such deviations have the potential to alter the Type I error rate of the test.

The random patterns and random shifts procedures performed similarly well when recreating the spatial patterns characterized by *Lagenifera*. This difference between *Poa* and *Lagenifera* is not surprising, given *Lagenifera* is characterized by much more fragmented patterns with smaller individual patch sizes, i.e., wrapping these patterns onto opposite sides of the area does not result in a significant alteration to the overall spatial structure. This suggests that species patterns that are characterized by a highly fragmented mosaic of patches are satisfactorily recreated by the random shifts test. However if the spatial pattern is dominated by one or

→

Fig. 3. (a) Fifty semivariograms generated by the random patterns and random shifts algorithms, based on the observed distribution map in part (b). The fine lines are the semivariograms for the 50 randomly generated patterns, and the thick line is the semivariogram for the observed pattern. Parts (b)–(f) show semivariograms for five observed distributions of *Poa sieberiana*. The spatial pattern on which the observed semivariogram is based is in the bottom right-hand corner of each graph. The symbols at each distance represent the average value over 50 randomly generated semivariograms. Open symbols indicate distances where the observed semivariogram value did not fall within the range of the 50 values generated by the respective randomization method.

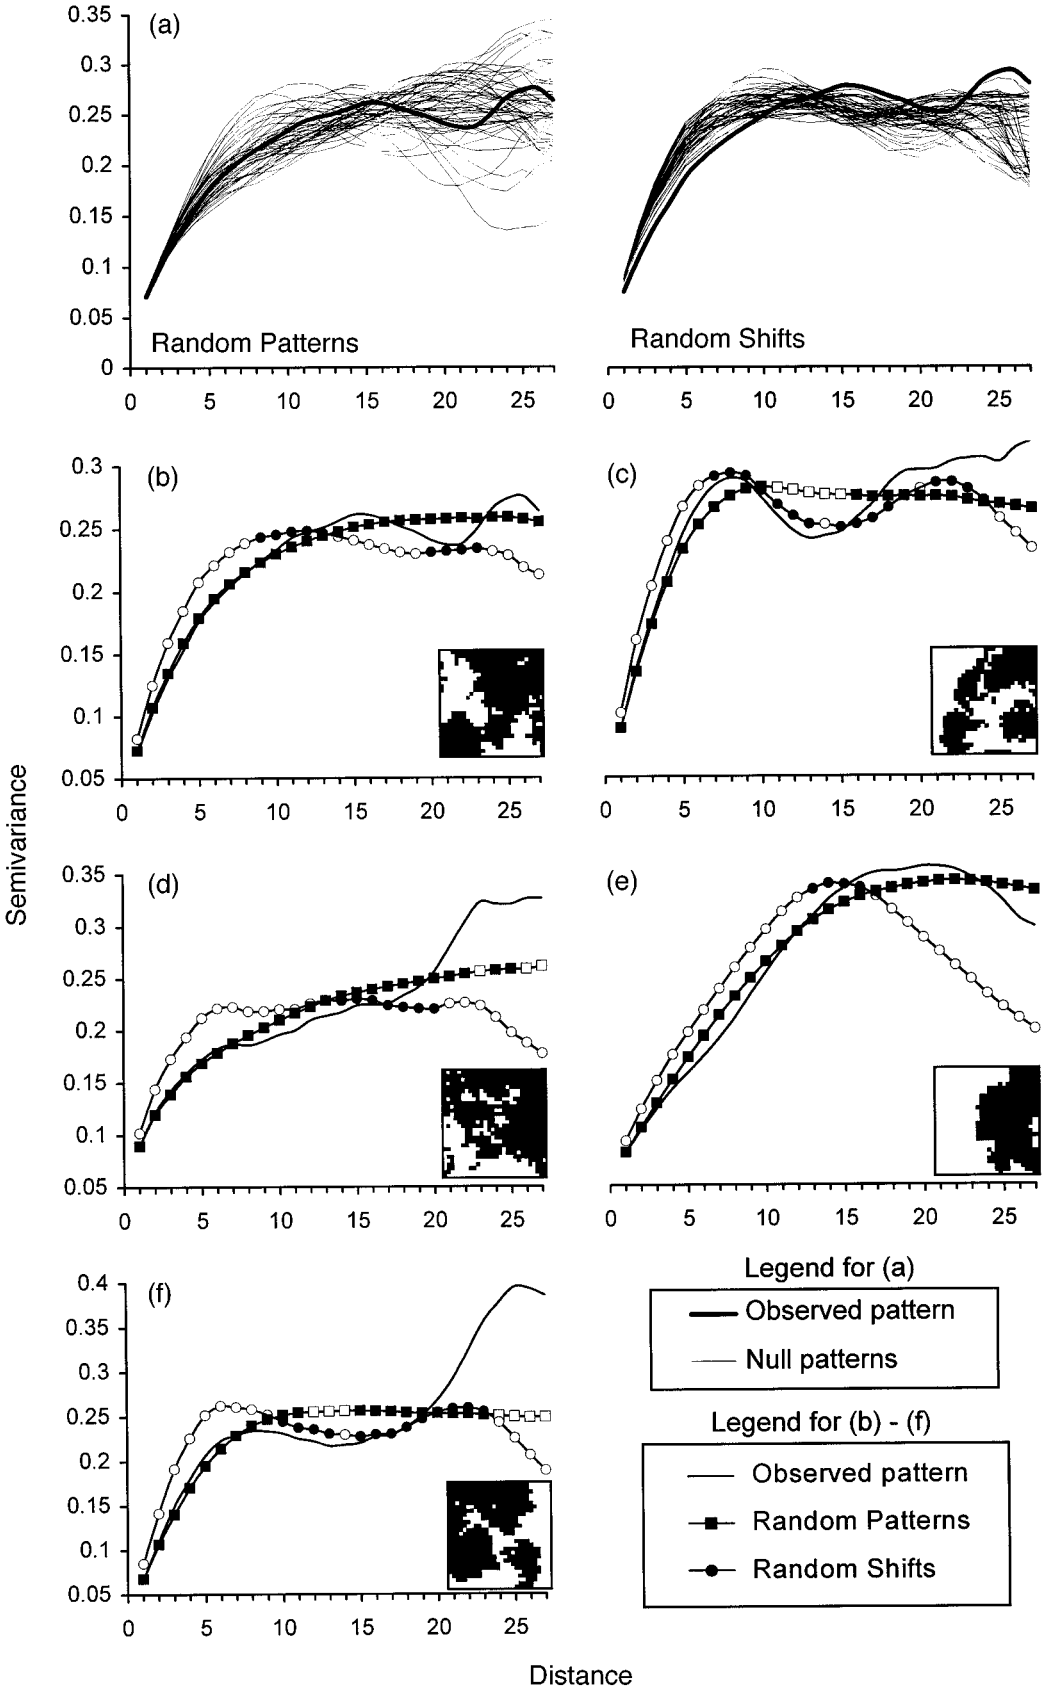
(a)
Random Patterns
Random Shifts
(b)
(c)
(d)
(e)
(f)
Semivariance
Distance
Legend for (a)
Observed pattern
Null patterns
Legend for (b) - (f)
Observed pattern
Random Patterns
Random Shifts

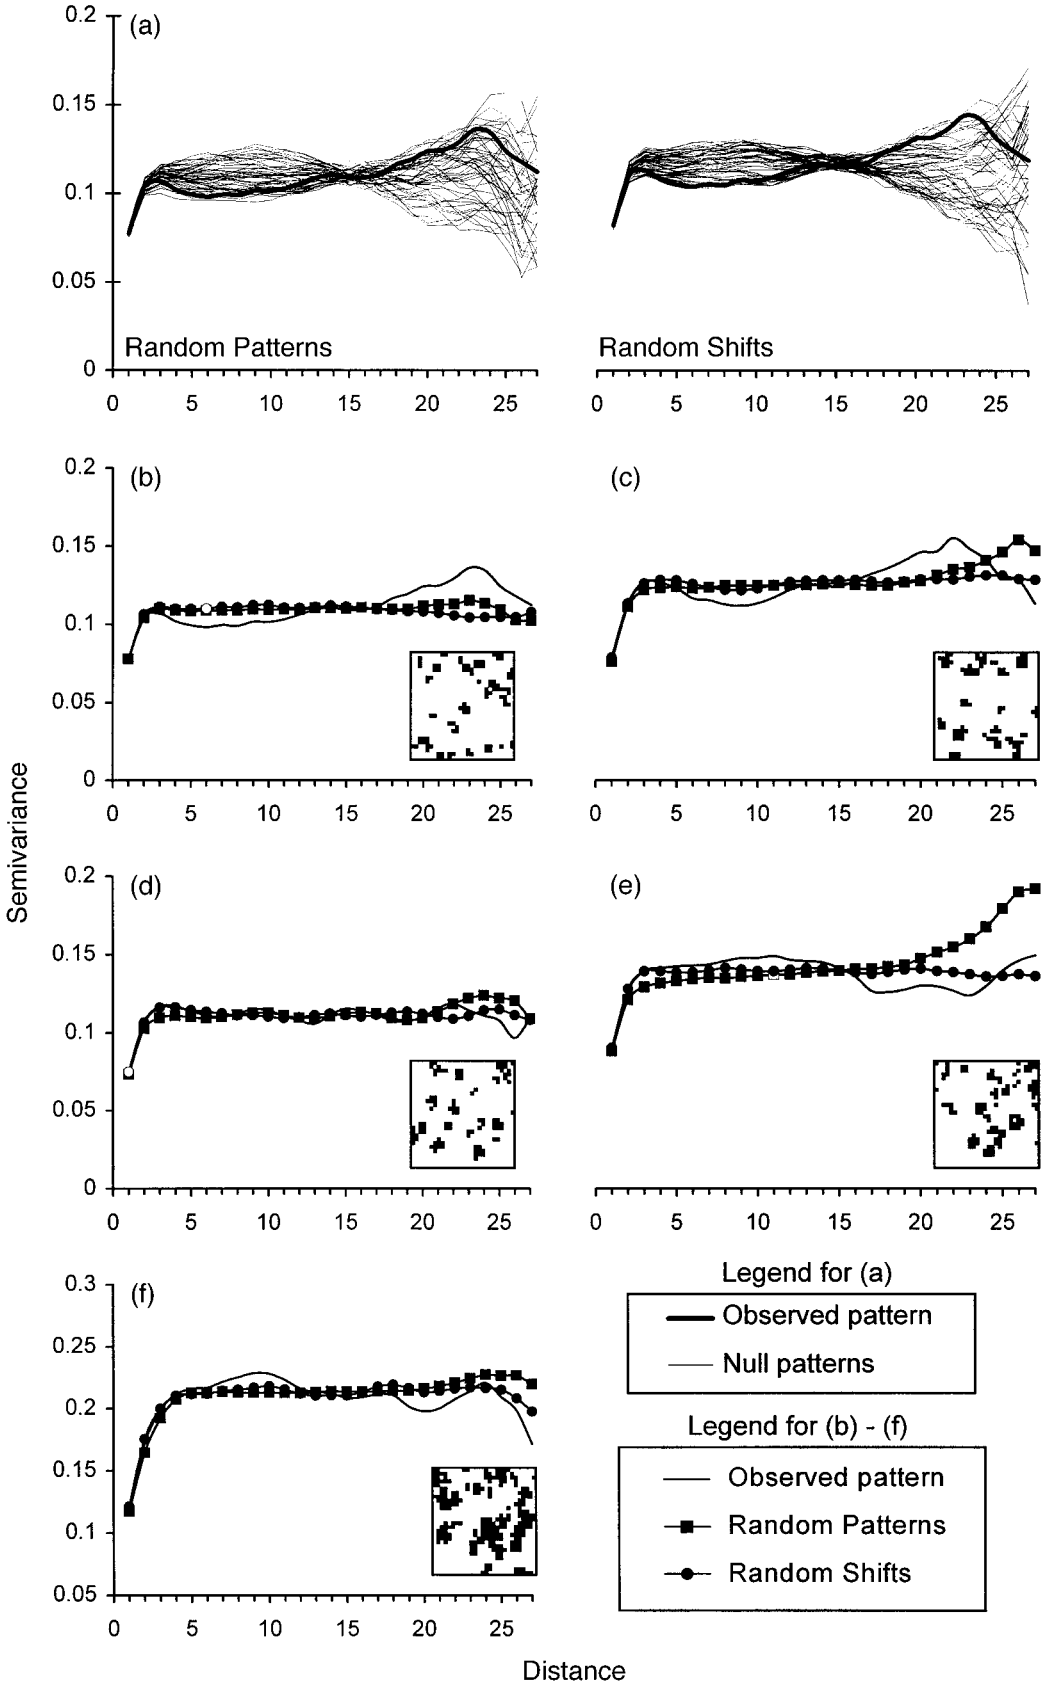


Fig. 4. Semivariograms for five observed distributions of *Lagenifera stipitata*. Format is as for Fig. 3.

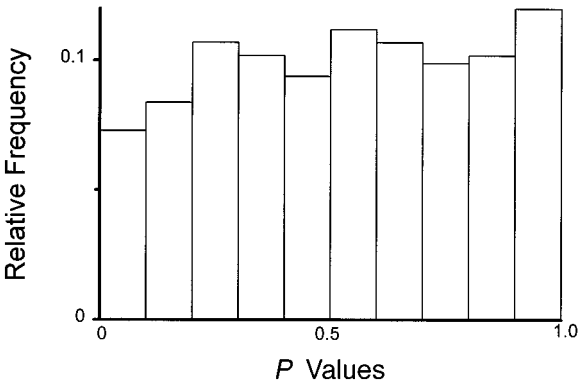


FIG. 5. Relative frequency histogram of the distribution of *P* values generated under the random patterns null model.

a few large patches, then a procedure such as the random patterns algorithm is the only way to ensure that the spatial characteristics at the whole-plot level are adequately recreated.

Despite the ability of the random patterns algorithm to simulate a wide range of observed patterns, there will always be some patterns with strong higher-order structures that it will not be able to satisfactorily recreate. For example, the random patterns semivariogram in Fig. 3c misses the dip at distances 10–15, corresponding to the particular ''concentric'' pattern of the patches evident in this map. One solution to this problem could be to add extra terms to the pattern generator that quantifies these higher-order structures. However it is likely that different sorts of higher-order structures would require different terms, e,g., gradients vs. ''patches within patches,'' and the extra computational overhead may make computing time prohibitive.

### *Confirmation of the distribution of the test statistic under the null model*

Although the random patterns procedure produces a very large number of distinct values of overlap, there is nevertheless still a finite number of possible values. I.e., the probability distribution of overlap is discrete. Such discreteness can mean the distribution of *P* values for the test departs from the desired rectangular distribution on [0, 1].

To check this, a procedure similar to Watkins and Wilson (1992) was used. For a plot of 28 × 28 cells (the dimensions of our experimental plots) two species were placed in the cells at random, each with a random frequency from 1 to 783 (one less than the total number of cells). This was our pseudo-observed distribution pattern which was subject to the random patterns test. This was repeated for 1000 pseudo-observed distributions to yield a probability distribution.

The frequency histogram of probability values obtained from the random data showed an approximately rectangular distribution, with some suggestion that the random patterns test is slightly conservative (Fig. 5). This phenomenon has also been observed by other workers (Watkins and Wilson 1992, Palmer and van der Maarel 1995) and can be attributed to the fact that with finite data there is a chance that the observed test statistic is occasionally exactly the same as the randomized. Because the probabilities from the randomization test are taken as the frequency of random test statistics that are equal to or more extreme than the observed, this method of calculation can result in an increase in the probabilities in the two tails to >1, resulting in a slightly conservative test. To correct for this conservatism Lancaster's ''median probability'' method could be used (Lancaster 1969). This method calculates *P* as [the probability of observing a random test statistic greater than that actually observed] + [0.5 × the probability of observing a random test statistic equal to that observed] (cf. Fig. 2).

TABLE 3. The eight species from the understorey community used in the analyses. Cover is the number of cells (1.4 cm × 1.4 cm) in which each species was recorded, divided by the total number of cells sampled and expressed as a percentage ($n$ = 56 448). Nomenclature follows Harden (1992).

| Species | Family | Cover |
|---|---|---|
| *Poa sieberiana* | Poaceae | 42.7 |
| *Lagenifera stipitata* | Asteraceae | 10.0 |
| *Microlaena stipoides* | Poaceae | 4.4 |
| *Viola hederacea* | Violaceae | 3.1 |
| *Geranium solanderi* | Geraniaceae | 2.8 |
| *Clematis aristata* | Ranunculaceae | 2.0 |
| *Asperula scoparia* | Rubiaceae | 1.5 |
| *Hydrocotyle peduncularis* | Apiaceae | 1.5 |

## APPLICATION OF THE RANDOM PATTERNS TEST

### *Site description and sampling methods*

The data used to illustrate the random patterns test were taken from part of a larger experiment investigating species coexistence in a community of herbaceous plants. The study site is located in the Brindabella Ranges near Canberra, Australia, at an altitude of ~1060 m. The canopy is dominated by a mixture of *Eucalyptus viminalis* and *E. fastigata.* The forest floor is characterized by a community of clonal, herbaceous plants. It is these species that are the focus of our study.

The experimental site from where the data were collected was 30 × 30 m square. Within this area are positioned, in a randomized split-block design, 36 pairs of plots. Each of these 72 plots is 40 × 40 cm square. Sampling was conducted in November 1994 and was achieved by subdividing each plot into a 28 × 28 grid of cells. Within each cell the shoot presence or absence of all living aboveground plant species (bryophytes + angiosperms) was recorded. It is this data which the random patterns test was applied. Twenty-two species were recorded overall, of which the eight most abundant were included for analysis (Table 3). Together these eight species comprised 94% of the total species records.

TABLE 4. Results of the 28 pairwise association tests among the eight species using the random patterns test. The *P* values are the probability of obtaining an overlap value as extreme as the one observed under the random patterns null model. The numbers in parentheses are the number of plots (out of 72) on which each analysis is based, i.e., the number of plots that contained both species. The sign indicates the direction of the association.

| Species | *Asperula scoparia* | *Clematis aristata* | *Geranium solanderi* | *Microlaena stipoides* | *Hydrocotyle peduncularis* | *Lagenifera stipitata* | *Poa sieberiana* |
|---|---|---|---|---|---|---|---|
| *Clematis aristata* | 0.172 − (35) | | | | | | |
| *Geranium solanderi* | 0.709 − (32) | 0.118 − (32) | | | | | |
| *Microlaena stipoides* | 0.126 − (30) | 0.623 + (38) | 0.312 + (32) | | | | |
| *Hydrocotyle peduncularis* | 0.060 + (24) | 0.854 + (25) | 0.284 − (17) | 0.666 − (23) | | | |
| *Lagenifera stipitata* | 0.213 − (48) | 0.078 − (56) | 0.023 − (44) | 0.812 − (47) | 0.106 − (30) | | |
| *Poa sieberiana* | 0.527 + (43) | 0.370 + (52) | 0.722 − (41) | 0.014 − (42) | 0.542 + (25) | 0.000 + (66) | |
| *Viola hederacea* | 0.148 + (33) | 0.342 − (38) | 0.502 − (34) | 0.034 + (33) | 0.054 + (20) | 0.298 − (52) | 0.062 + (50) |

Note that the example used to illustrate the random patterns test (Fig. 1) was based on a single 40 cm square plot, with the test statistic being the overlap of the species within this particular plot. However, for the community analysis presented below the overall-plots association was used as the test statistic. The value of this statistic for the observed data was calculated as the sum of the overlap values for each plot in which both species occurred. The null distribution of this statistic was obtained by applying the random patterns algorithm to each plot in turn, and summing the random overlap values for each plot to obtain, for a single randomization, an overall-plots overlap under the null model. To generate the null distribution, $n = 1000$ randomizations were performed. By pooling the information from a number of plots in this way we ensured that any deviations from randomness reflected consistent trend across the whole sampling area, and not just the particular characteristics of a single 40 cm square plot.

### *Results and discussion*

The results of applying the random patterns test to all pairwise combinations of the eight species are presented in Table 4.

With 28 pairwise comparisons some tests would be expected to be significant by chance alone, without any genuine associations between species. To guard against this possibility a Bonferroni correction can be applied, in which case an individual *P* value would be judged significant at the 5% level if it is $<0.05/x$, where $x$ is the number of tests performed. Applying this correction to our results means that for any individual test to be significant it must have probability value of $<0.05/28 = 0.0018$. Under this criterion there was only one association out of the 28 comparisons that was statistically significant, a positive association between *Poa* and *Lagenifera* ($P < 0.000$). From this result we can be confident that associations between species do exist based on our results, and that at least *Poa* and *Lagenifera* are significantly associated. However, this procedure is highly conservative, and it is a reasonable expectation that some of the other species are genuinely associated as well, for example, *Poa* and *Microlaena*, and *Lagenifera* and *Geranium*, have low *P* values, which are suggestive of significance but are not judged significant by the Bonferroni criterion.

It has been pointed out by a number of authors that observing a significant association between two species does not logically imply any specific ecological process (e.g., Schluter 1984). Analyses such as those presented in Table 4 should therefore be used in relation to other information, to be followed by experimentation. This is the approach we have taken in our study, where a number of experiments have been established to enable us to investigate in greater detail the association between *Poa Lagenifera.*

A further criticism of the multiple pairwise approach is that the degree of association might depend on what other species are present. For example, species 1 and 2 might be negatively associated, but only when species 3 is present. To overcome such problems a multi-species contingency table analysis can be used (Dale et al. 1991). We argue that such criticisms do not necessarily invalidate the use of pairwise association tests, but rather reflect two different ways of analyzing and interpreting patterns of association. The first considers all other species in the community to be constant, and concentrates only on the pairwise association at hand. The second recognizes that the species are embedded within a multi-species community, and therefore can be used to investigate multi-species association. When used together, and providing the above caveats are kept

in mind, then the two methods can be considered complimentary, and hence have the potential to provide us with more information than if either method were used alone.

In this paper we have concentrated only on the pairwise association analyses. The reasons for doing so were twofold. First, we are currently implementing the random patterns test as part of a larger study investigating the multi-species association patterns at a range of both spatial and temporal scales (S. H. Roxburgh and P. Chesson, *unpublished data*). Second, the main aim of this paper is to introduce the random patterns test, and the less complex pairwise association analysis facilitated a clearer description and demonstration of the method.

### *Comparison with the chi-squared test of association*

To further highlight the dangers of applying traditional statistical techniques when the data are spatially autocorrelated, the results of the standard contingency table analysis were applied to the same 28 pairs of species. These analyses detected nine significant associations (results not presented), compared with the single significant association detected with the random patterns test. Clearly, indiscriminate use of the traditional techniques, without first confirming within-species randomness, has the potential to seriously mislead.

## Conclusions

The underlying assumption of the random patterns test is that the spatial patterns are a reflection of the growth process of each species (or migration etc.), and by including the characteristics of the spatial pattern into the null model we are mimicking, in the randomizations, the growth processes of the species within each plot. This ensures that the patterns constructed under the null model are biologically sensible, while at the same time taking into account the detrimental effects of spatial autocorrelation.

The random patterns method was applied in the context of a test for pairwise association, however its utility is potentially far broader. For example, tests of multi-species association are similarly affected by spatial autocorrelation, and the random patterns method is easily extended to include these (Roxburgh and Chesson, unpublished data). More generally, the random patterns approach can be used to test a wide range of hypotheses involving the analysis of spatial data. For example, tests for niche limitation (Watkins and Wilson 1992, Palmer and van der Maarel 1995), and guild proportionality (Wilson and Roxburgh 1994) are also potentially affected by spatial autocorrelation, and both are amenable to the random patterns approach.

## Acknowledgments

We thank Philip Dixon, J. Bastow Wilson, and two anonymous reviewers for comments.

## Literature Cited

Burrough, P. A. 1987. Spatial aspects of ecological data. Pages 213–251 *in* R. H. Jongman, C. J. F. ter Braak, and O. F. R. van Tongeren, editors. Data analysis in community and landscape ecology. Center for Agricultural Publishing and Documentation (PUDOC), Wageningen, The Netherlands.

Cressie, N. A. 1993. Statistics for spatial data. John Wiley and Sons, New York, New York, USA.

Dale, M. R. T. 1977. Graph theoretical analysis of the phytosociological structure of plant communities: the theoretical basis. Vegetatio **34**:137–154.

Dale, M. R. T., D. J. Blundon, D. A. MacIsaac, and A. G. Thomas. 1991. Multiple species effects and spatial autocorrelation in detecting species associations. Journal of Vegetation Science **2**:635–642.

Edgington, E. S. 1995. Randomization tests. Third edition. Marcel Dekker, New York, New York, USA.

Goodchild, M. F. 1986. Spatial autocorrelation. Geo Books, Norwich, UK.

Greig-Smith, P. 1983. Quantitative plant ecology. Third edition. Blackwell Scientific, Oxford, UK.

Harden, G. J. 1992. Flora of New South Wales. New South Wales University Press, Kensington, New South Wales, Australia.

Hurlbert, S. H. 1984. Pseudoreplication and the design of ecological field experiments. Ecological Monographs **54**:187–211.

Journal, A. J., and J. C. Huijbregts. 1978. Mining geostatistics. Academic Press, London, UK.

Lancaster, H. O. 1969. The chi-squared distribution. John Wiley and Sons, New York, New York, USA.

Legendre, P. 1993. Spatial autocorrelation - trouble or new paradigm. Ecology **74**:1659–1673.

Ludwig, J. A., and J. F. Reynolds. 1988. Statistical Ecology. John Wiley and Sons, New York, New York, USA..

Mangel, M., and F. R. Adler. 1994. Construction of multidimensional clustered patterns. Ecology **75**:1289–1298.

Manly, B. F. J. 1991. Randomization and Monte-Carlo methods in biology. Chapman and Hall, London, UK.

McCulloch, C. E. 1985. Variance tests for species association. Ecology **66**:1676–1681.

McElhany, P., L. A. Real, and A. G. Power. 1995. Vector preference and disease dynamics: a study of barley yellow dwarf virus. Ecology **76**:444–457.

Myster, W., and S. T. A. Pickett. 1992. Dynamics of associations between plants in ten old fields during 31 years of succession. Journal of Ecology **80**:291–302.

Noreen, E. W. 1989. Computer-intensive methods for testing hypotheses: an introduction. John Wiley and Sons, New York, New York, USA.

O'Connor, I., and L. W. Aarssen. 1987. Species association patterns in abandoned sand quarries. Vegetatio **73**:101–109.

Palmer, M. W. 1992. The coexistence of species in fractal landscapes. American Naturalist **139**:375–397.

Palmer, M. W., and E. van der Maarel. 1995. Variance in species richness, species association, and niche limitation. Oikos **73**: 203–213.

Pielou, E. C. 1977. Mathematical ecology. John Wiley and Sons, New York, New York, USA.

Real, L. A., and P. McElhany. 1996. Spatial pattern and process in plant-pathogen interactions. Ecology **77**:1011–1025.

Rejmánek, M., and J. Lepš. 1996. Negative associations can reveal interspecific competition and reversal of competitive hierarchies during succession. Oikos **76**:161–168.

Rossi, R. E., D. J. Mulla, A. G. Journel, and E. H. Franz. 1992. Geostatistical tools for modeling and interpreting ecological spatial dependence. Ecological Monographs **62**:277–314.

Schluter, D. 1984. A variance test for detecting species associations, with some applications. Ecology **65**:998–1005.

Tavaré, S., and P. M. E. Altham. 1983. Serial dependence of observations leading to contingency tables, and corrections to chi-squared statistics. Biometrika **70**:139–144.
Wackernagel, H. 1995. Multivariate geostatistics. Springer-Verlag, Berlin, Germany.
Watkins, A. J., and J. B. Wilson. 1992. Fine-scale community structure of lawns. Journal of Ecology **80**:15–24.
Whittaker, R. H., and W. A. Niering. 1975. Vegetation of the Santa Catalina Mountains, Arizona. V. Biomass. Production and diversity along the elevation gradient. Ecology **56**:771–790.
Wilson, J. B. 1995. Variance in species richness, niche limitation, and vindication of patch models. Oikos **73**:277–279.
Wilson, J. B., and S. H. Roxburgh. 1994. A demonstration of guild-based assembly rules for a plant community, and determination of intrinsic guilds. Oikos **69**:267–276.

## APPENDIX

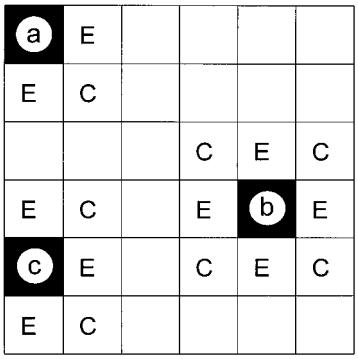


FIG. A1. The Edge (*E*) and Corner (*C*) definitions used to quantify the within-species spatial patterns for the random patterns algorithm. There are three possible positions within the map: (a) the map-corner position, (b) the map-interior position, (c) the map-edge position. The number of adjoining cells for each statistic depends upon the position of the target cell within the map. E.g., map-interior target cells have four possible edge cells and four possible corner cells, and map-edge target cells have three possible edge cells and two corner cells.

### *Statistics used to quantify the within-species spatial patterns*

Within-species spatial patterns were quantified using four statistics: edge (*E*) and corner (*C*) contacts, and open (*O*) and solid (*S*) areas. These statistics are defined as follows.

Edge contacts (*E*): The edge cells are defined relative to the rook moves in chess (Fig. A1). For each target cell that is occupied by a species, the total number of adjacent edge cells that are also occupied by the same species are counted and the result expressed as a proportion of the total possible edge contacts for that cell. The total number of possible edge cells depends on the location of the target cell within the overall map (see Fig. A1). For all occupied cells these scores are summed, and for convenience the average of these values can be calculated. For example the value 0.83 for *Poa* in Fig. 1a indicates that where a cell contains *Poa*, on average 83% of the surrounding edge cells also contain *Poa*.

Corner contacts (*C*): This is a statistic similar to the edge contacts, but calculated as for the equivalent bishop moves in chess (see Fig. A1).

Open areas (*O*): This is the total number of times a target cell without the species is totally surrounded by empty cells. It gives a measure of how open the bare patches are in the pattern. The cells at the corner and at the edges of the pattern were treated similarly to the calculation of the edge (*E*) and corner (*C*) contacts.

Solid areas (*S*): This is the converse to the open areas measure and gives a measure of how extensive the solid patches are in the pattern. That is, in target cells where the species is present, *S* is the total number of times all possible surrounding cells are also occupied by that species.

### *Techniques to increase the performance of the random patterns algorithm*

Implementing the random patterns procedure as we have described them is extremely computer intensive. For example, Fig. 1a shows 506 cells occupied by *Poa.* To recreate the patterns for *Poa* in Fig. 1c from initial random assignments such as those in Fig. 1b requires ~5500 individual cell swapping events. To increase the efficiency of the algorithm the randomization procedure was given a ''head start'' as follows. Instead of the initial randomization allocating all of the individual cells independently as in Fig. 1b, the overall area was divided into blocks, and these blocks were randomly shuffled. For the maps in Fig. 1 the area was divided into a $4 \times 4$ grid of blocks, each with dimensions $7 \times 7$ cells. For a given pattern, the exact dimensions of the blocks and their total number will be a function of the dimensions of the plots being analyzed. Each block was then subject to a random rotation/reflection, and then these 16 blocks were assigned to random positions within the plot. From this initial ''block'' randomization the random patterns procedure was then implemented, shuffling the $7 \times 7$ blocks of cells about.

This procedure is repeated until the fit between the observed and randomized pattern falls below $\alpha$, or does not improve after 100 swapping attempts. If 100 swapping attempts are attained without a satisfactory fit, further ''fine-tuning'' of the random pattern is required. This is achieved by fixing the 16 blocks in position, and then shuffling the single cells until the desired function value of $\alpha$ was achieved. That is, the randomization procedure comprised two stages, firstly blocks at the scale of $7 \times 7$ cells, followed by individual cells. Performing this two-stage procedure has two main advantages. Firstly, the algorithm runs much faster because the initial randomization based on the shuffling of blocks requires less ''re-shuffling'' to create a spatial pattern with the same characteristics as the original. For example the number of individual cell-swapping events required to recreate the *Poa* distribution using this modified method reduced from 5500 to ~950 events. Note that splitting the area into 16 blocks still ensures adequate randomization over the plot, thus avoiding the problems inherent with the rotation/reflection method. Secondly, we found that starting from a completely random pattern occasionally failed to converge to a satisfactory solution for observed patterns which had e.g., a small number of discrete patches. By retaining a certain amount of spatial structure after the initial randomization prevented this problem of nonconvergence.