Ecology, 87(8), 2006, pp. 1925–1931

STATISTICAL INFERENCE USING THE *G* OR *K* POINT PATTERN SPATIAL STATISTICS

N. Bert Loosmore[1,3] and E. David Ford[1,2]

[1]*Quantitative Ecology and Resource Management, Box 352182, University of Washington, Seattle, Washington 98195-2182 USA*
[2]*College of Forest Resources, Box 325100, University of Washington, Seattle, Washington 98195-2100 USA*

Abstract. Spatial point pattern analysis provides a statistical method to compare an observed spatial pattern against a hypothesized spatial process model. The *G* statistic, which considers the distribution of nearest neighbor distances, and the *K* statistic, which evaluates the distribution of all neighbor distances, are commonly used in such analyses. One method of employing these statistics involves building a simulation envelope from the result of many simulated patterns of the hypothesized model. Specifically, a simulation envelope is created by calculating, at every distance, the minimum and maximum results computed across the simulated patterns. A statistical test is performed by evaluating where the results from an observed pattern fall with respect to the simulation envelope. However, this method, which differs from P. Diggle's suggested approach, is invalid for inference because it violates the assumptions of Monte Carlo methods and results in incorrect type I error rate performance. Similarly, using the simulation envelope to estimate the range of distances over which an observed pattern deviates from the hypothesized model is also suspect. The technical details of why the simulation envelope provides incorrect type I error rate performance are described. A valid test is then proposed, and details about how the number of simulated patterns impacts the statistical significance are explained. Finally, an example of using the proposed test within an exploratory data analysis framework is provided.

Key words: *G, K, and L statistic; Monte Carlo; simulation envelope; spatial point patterns.*

Introduction

A point pattern consists of the spatial configuration of observed events within a specified area, such as Fig. 1, where the points represent the location of trees in a 1-ha plot of the Wind River Canopy Crane Research Facility (WRCCRF). Spatial point pattern statistics utilize the distribution of distances between events to make inference about the spatial arrangement. Specifically, the *G* statistic (Eq. 1) relies upon the distribution of nearest-neighbor distances and the *K* (or *L*) statistic (Eq. 2 without edge correction) utilizes the distribution of all interpoint distances (Diggle 2003: Chapters 2 and 4). In these equations, t is the distance, n is the number of points within the given pattern, d_i is the distance from point i to its nearest neighbor, d_{ij} is the distance from point i to point j, and $I(\ )$ is an indicator function that equals 1 if the argument is true and 0 otherwise. Both the *G* and *K* (or *L*) are cumulative and calculated as a function of increasing distance between events:

$$G(t) \approx \frac{1}{n}\sum_{i=1}^{n}[I(d_i < t)] \tag{1}$$

$$K(t) \approx \frac{1}{n(n-1)}\sum_{i=1}^{n}\sum_{j=1, j\neq i}^{n}[I(d_{ij} < t)]. \tag{2}$$

These statistics can be used to test whether an observed point pattern has similar characteristics to a specified spatial model, such as complete spatial randomness (CSR), and to evaluate whether a pattern exhibits aggregation or inhibition. Additionally, it has been considered that these statistics can provide information about scale, i.e., the distances at which an observed pattern deviates from the specified null-process model.

A simulation envelope is created by using at every distance the minimum and maximum results for a chosen statistic (*G* or *K*), calculated from a number of Monte Carlo (MC) simulated patterns. Patterns are simulated from the hypothesized spatial model. Observed results, also calculated from the chosen statistic, are plotted on top of this envelope and compared. The

Manuscript received 29 September 2005; revised 24 February 2006; accepted 9 March 2006. Corresponding Editor: N. J. Gotelli.
[3] E-mail: nhl@u.washington.edu

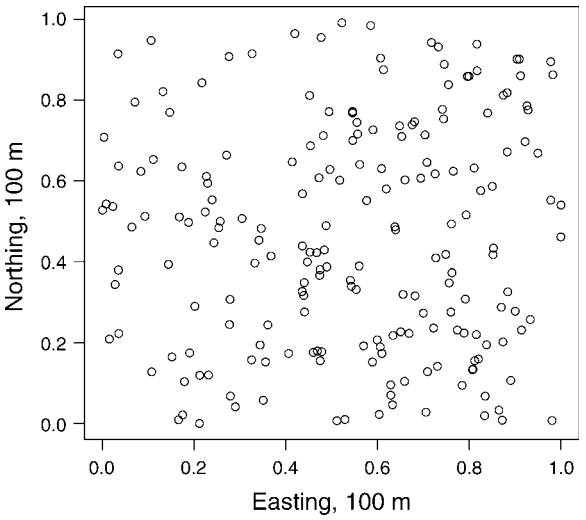

FIG. 1. Plot of locations of all trees <6 m tall, indicated by the circles, from a 1-ha (100 × 100 m) plot of the Wind River Canopy Crane Research Facility (WRCCRF) stem map, scaled to a unit square (1.0 corresponds to 100 m).

observed process is deemed different than the simulated process if its results lie outside the envelope at any distance. Kenkel (1988) pioneered the use of the simulation envelope with the K statistic to make inferences about competition occurring within a jack pine stand, and this method of using spatial point pattern statistics is now common in the literature. As Table 1 shows, variation exists within the statistics used, the number of simulated patterns from which the envelope was developed, and the significance of the confidence interval (CI) assigned to the envelope limits. Variation even exists in the relationship between the number of simulated patterns and the CI. The studies listed comprise only a subset of recently published articles employing these methods.

Unfortunately, this approach is invalid for testing an observed point pattern against a specified spatial model because the type I error rate is greatly underestimated. For similar reasons, determining the distances at which the observed pattern differs from the spatial model based on the distances at which the results of the observed pattern exceed the simulation envelope is also suspect. Our first objective in this paper is to detail exactly why the simulation envelope approach is inappropriate. We then propose a goodness-of-fit (GoF) test which provides expected type I error rates, and explain how the number of simulated patterns should be chosen and its relationship to the significance of the test. Finally, we conclude with a simple example illustrating our method.

METHODS

Failure of the simulation envelope

The perceived, although inaccurate, type I error rate of a simulation envelope based statistical test is $\alpha = 1/s$, where $s - 1$ is the number of patterns used to construct the envelope, and the additional 1 in the denominator accounts for the observed pattern being tested. This result would be valid only for a one-sided test where the limit of the envelope was comprised of the results from a single pattern. However, researchers commonly apply the simulation envelope in a two-sided fashion, rejecting H_0 when the results from the observed pattern fall either above or below the envelope. Furthermore, the envelope limits are typically comprised from the results of many simulated spatial point patterns, each contributing to the envelope over different distances. Freeman and Ford (2002: their Fig. 2) demonstrate this by displaying the progression of a simulation envelope being built. Similarly, our Fig. 2 shows an example of the limits of a simulation envelope for the G statistic built from 99 CSR patterns. In this example, 74 of the patterns actually comprise the envelope at some distance. The value $\hat{G} - \overline{G}$, where $\hat{G}$ represents either the upper or lower envelope result and $\overline{G}$ represents the mean result calculated across all simulations, is plotted to enhance the details of the envelope and center the results around $y = 0$.

A type I error is defined, in general, when H_0 is incorrectly rejected even though it is true. For the envelope shown in Fig. 2, a type I error would occur whenever a new CSR pattern fell outside the envelope. In this case, the probability that a CSR pattern would exceed the displayed envelope is $74/99 \approx 0.74$; much higher than the expected type I error rate of $1/s = 0.01$. The weaker than expected statistical performance is the result of many tests, one at each different distance, being performed concurrently. This simultaneous inference (e.g., Gotelli and Ellison 2004: Chapter 10)

TABLE 1. Recent ecological examples in which the G and/or K statistics were used to analyze observed spatial point patterns with the statistic(s) employed, the number of patterns comprising the simulation envelope, and the perceived significance of the "confidence interval" (CI) created by the simulation envelope.

Example	Statistic(s)	No. patterns	CI (%)
Batista and Maguire (1998)	G, K	19	95
Dolezal, Stastna, Hara, and Srutek (2004)	K	99	99
Druckenbrod, Shugart, and Davies (2005)	L	1000	95
Freeman and Ford (2002)	G, K	99	99
Grassi et al. (2004)	K	99	99
Hirayama and Sakimoto (2003)	K	19, 99	95, 99
Martens et al. (1997)	L	99	95
Moeur (1997)	G, K	200	90
North et al. (2004)	K	100	95
Parish, Antos, and Fortin (1999)	G, K	99	99
Rigg (2005)	K	99	99
Salvador-Van Eysenrode et al. (2000)	G, K	1000	95
Srutek, Dolezal, and Hara (2002)	L	99	95
Tirado and Pugnaire (2003)	K	1000	95
Wiegand and Moloney (2004)	L	99	99

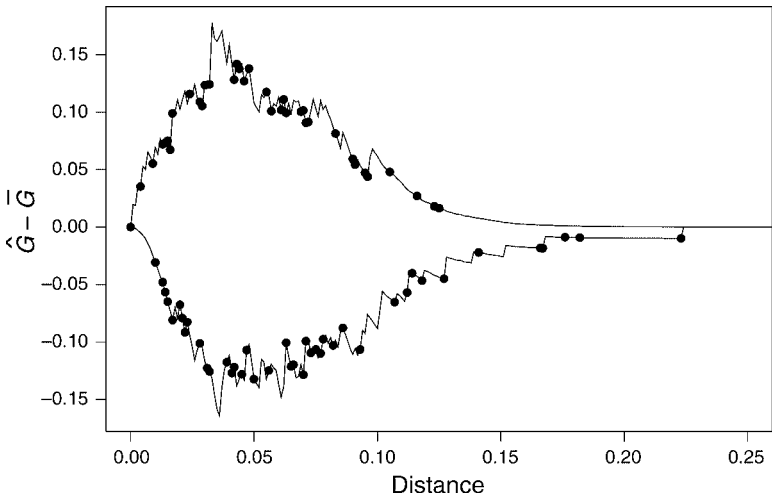

FIG. 2. Plot of $\hat{G} - \overline{G}$ against distance for a simulation envelope built from $s = 99$ simulations of a complete spatial randomness (CSR) model for $\lambda = 100$. $\hat{G}$ represents either the upper or lower envelope result, and $\overline{G}$ represents the mean result calculated across all simulations. The circles represent transitions when a new pattern begins to comprise that section of the envelope. In this example, 74 of the 99 simulated patterns contribute to the envelope at some distance. Distance is a unit scale, with 0.10 corresponding to 10 m.

yields an underestimated type I error rate and therefore equating the upper and lower bounds of the simulation envelope with a $(1 - \alpha)\%$ CI is invalid. Although multiple testing problems are often solved using a Bonferroni correction, it is not appropriate here because of both the correlation between results at consecutive distances which violates the assumption of independence, as well as the large number of distances being simultaneously evaluated.

Further, the method of inferring scale, i.e., the range of distances over which the observed pattern differs from the hypothesized model, based on where the observed results exceed the envelope is also invalid. Not only is the type I error rate incorrect, but because of the cumulative nature of the G and K statistics, results at any distance reflect both the instantaneous value at that distance as well as the combined results from all smaller distances. Results for an observed pattern could therefore lie outside the envelope at a distance where the instantaneous value was not different than the specified model.

A valid goodness-of-fit test

As a replacement for the simulation envelope, we propose a goodness-of-fit (GoF) test, based on the approach of Diggle (2003: Chapter 2) and Cressie (1991: Section 8.4). Under this GoF test, both the observed pattern and each of the MC simulated patterns are each reduced to a single summary test statistic, as

$$u_i = \sum_{t_k=t_{\min}}^{t_{\max}} [\hat{H}_i(t_k) - \overline{H}_i(t_k)]^2 \delta t_k \qquad (3)$$

where t_k is distance, $t_{\min}$ and $t_{\max}$ are the lower and upper limits of the summation in terms of distance, $\hat{H}_i(t_k)$ is the empirical result for pattern i for the test statistic of interest (G or K), $\overline{H}_i(t_k)$ is the mean result computed for all patterns except for i, and $\delta t_k = (t_{k+1} - t_k)$ is the width of the distance interval.

$\overline{H}_i(t)$ is used to reduce any bias that may be present from edge effects or from estimating the number of points within the process (Diggle 2003):

$$\overline{H}_i(t) = \frac{1}{s-1}\left[\sum_{j=1, j\neq i}^{s} \hat{H}_j(t)\right]. \qquad (4)$$

While Eq. 3 might utilize the theoretical value of the process, $H(t)$ in place of $\overline{H}(t)$, the theoretical value is rarely known. Additionally, based on the weak law of large numbers (Casella and Berger 2002: Chapter 5), $\lim_{s\to\infty} \overline{H}_i(t) \to H(t)$. Hence this summary statistic value (u_i) represents the total squared deviation between the observed pattern and the theoretical result across the distances of interest.

A vector that describes the specific distances ($t_k \in t_1 = t_{\min}, t_2, \ldots, t_n = t_{\max}$) at which the observed and simulated patterns are to be evaluated is needed to compute each u_i. To ensure all changes in $\hat{G}$ or $\hat{K}$ are properly detected, we recommend using the empirical distance list, i.e., the vector of all distances where the chosen statistic changes, as determined from pre-sampling the simulated and observed patterns.

A statistical test (i.e., rejecting or failing to reject H_0) can be performed by ranking the summary statistic value for the observed pattern (u_1) within the results from the $s - 1$ simulated patterns, u_i, $i = 2, \ldots, s$. To be clear, the explicit null hypothesis being tested here is not that the observed pattern was generated by the null spatial

process. Instead H_0 is that the summary statistic value, u_1, calculated for the observed pattern for the chosen statistic (G or K) over the range of distances $\{t: t_{\min} \leq t \leq t_{\max}\}$ specified is not different than those calculated for random instances of the hypothesized spatial process model.

The observed P value of this hypothesis test ($\hat{p}$) is calculated as

$$\hat{p} = 1 - \frac{\text{rank}[u_1] - 1}{s} = 1 - \frac{\sum_j I[u_1 > u_j]}{s} \qquad j = 2, \ldots, s \tag{5}$$

where $I[u_1 > u_j]$ is an indicator function that equals 1 if u_1 is greater than the given u_j and 0 otherwise. For example, if the summary statistic computed for the observed pattern were larger than that of each of the $s - 1 = 99$ simulated patterns, then $\hat{p} = 1 - 99/100 = 0.01$.

To summarize, the proposed GoF test converts the multivariate result (i.e., the calculated statistic as a function of distance) for an observed pattern into a univariate summary test statistic. MC theory can then be used to accept or reject H_0, maintaining the expected type I error rate.

Characterizing uncertainty in $\hat{p}$

Since this test relies on simulated patterns to create the null distribution of u_i values, some amount of uncertainty exists within the realized P value ($\hat{p}$) of the observed pattern when tested against a specified null process model. Increasing the number of simulated patterns will reduce this uncertainty, but the true P value (p) for an observed pattern is ultimately unknowable. Hence, it is worth understanding how the number of simulations used to build the null distribution impacts the amount of variation in $\hat{p}$. Marriott (1979) referred to this issue as a blurred critical region. We take a different approach using distribution theory to determine the variation in $\hat{p}$ as a function of s.

To derive this variation, let Y represent the summary statistic value from the data (u_1) and, similarly, let X_j represent the u_j value from the $j = 1, \ldots, s - 1$ simulations of the null process. The realized P value from Eq. 5 can be rewritten as

$$\hat{p} = 1 - \frac{\sum_j I(Y > X_j)}{s}. \tag{6}$$

Under H_0, $\hat{p}$ has a discrete uniform distribution on $1/s, 2/s, \ldots, s/s$. Assuming Y comes from H_0, then by definition, the true P value is $p = \Pr(X > Y \mid Y)$. Hence, each of the $I(Y > X)$ is distributed Bernoulli with probability $1 - p$. The expected value of $\hat{p}$ is

$$\mathrm{E}(\hat{p}) = \mathrm{E}\left[1 - \frac{\sum_j I(Y > X_j)}{s}\right] = 1 - \frac{(s-1)(1-p)}{s} \approx p \tag{7}$$

showing that the test is unbiased for large s. Similarly, the variance of $\hat{p}$ (denoted as σ_p^2) is

$$\sigma_p^2 = \mathrm{Var}\left[1 - \frac{\sum_j I(Y > X_j)}{s}\right] = \frac{1}{s^2}(s-1)p(1-p) \approx \frac{p(1-p)}{s} \tag{8}$$

for large s. We can therefore approximate the distribution of $\hat{p}$ using a binomial distribution.

Estimating σ_p^2 based on Eq. 8 requires a value for p. Given a hypothetical value of $\hat{p}$, the theoretical variation can be determined. For example, we might be interested in knowing the theoretical variation of $\hat{p}$ for a pattern that might fail a hypothesis test, i.e., a pattern with a $\hat{p}$ near 0.10. As such, we can set $p = 0.10$ and apply Eq. 8 to determine the amount of variation as a function of the number of simulated patterns. This is useful because the choice of the number of simulated patterns needs to be made before the GoF test is performed. Then, after a test is complete, the determined value of $\hat{p}$ can be used to directly estimate the specific variation of the given pattern and test.

To select the size of the null distribution, one must choose an acceptable amount of uncertainty for the observed P value. Given that binomial distributions converge to normal distributions for a large number of trials, a 95% confidence interval (CI) for the true P value of a pattern (p) can be written as

$$\{\hat{p} - 1.96 \times \sigma_p \leq p \leq \hat{p} + 1.96 \times \sigma_p\}. \tag{9}$$

The CI width determines the value of σ_p^2, which in turn is used to choose the number of simulated patterns. For example, if a CI width of 0.04 were acceptable, this would require a value of $\sigma_p \approx 0.01$. Using this value and an estimated $p = 0.10$ within Eq. 8 results in $s \approx 1000$.

A Simple Model of Establishment

Wiegand et al. (2003) describe how patterns in data can aid model development, and provide an example for spatial point pattern analysis. While their example incorrectly uses the simulation envelope approach, we agree with their four steps of a data based approach to model fitting: (1) aggregation of biological information, (2) determination of parameter values, (3) systematic comparison between observed patterns and those predicted by the model, and (4) secondary predictions. Our GoF test is a necessary component of step 3 when evaluating spatial point patterns.

Studies on a number of old-growth conifer forests have indicated that the distribution of all trees is clumped, and since these plots are dominated by smaller trees, researchers have inferred an aggregated pattern of establishment (e.g., Moeur 1993, 1997, Van Pelt and Franklin 2000, Freeman and Ford 2002, Harris 2004). Most studies have evaluated the distribution of tree locations using the simulation envelope approach and against an H_0 of CSR; we instead evaluate it using the

GoF test and against a more informative null spatial process model. Our observed data consist of the locations of all trees <6 m in height (Fig. 1) within a 1-ha plot of the WRCCRF, and such trees were chosen as these locations demonstrate the most recent aspects of the establishment process. A visual assessment suggests some amount of clustering in the locations of the trees.

For our null spatial process model, we chose a Poisson cluster model, wherein ρ parents are first randomly located on the plot using a random process and then $X \sim$ Poiss(μ) children are placed near each parent. This model may be appropriate given the patterns of seed distribution and seedling establishment. The expected point intensity λ is calculated as $\lambda = \rho\mu$. The spatial positions of children relative to their parents were chosen to be identically and independently distributed according to the bivariate normal distribution:

$$h(x, y) = \frac{1}{2\pi\sigma^2}\exp\left[-\frac{x^2 + y^2}{2\sigma^2}\right] \quad (10)$$

so the amount of aggregation will depend on the σ^2 term. Additionally, x, y represent spatial coordinates and the distribution of children points is isotropic and stationary.

To implement this model, estimates of parameters ρ, μ, and σ^2 were needed; λ was estimated directly from the number of points in the observed plot. Once a value for ρ was available, μ could then be determined directly as $\mu = \lambda/\rho$, but in the absence of guidance based on ecological theory, how could ρ be determined? Hence, the challenge here was how to determine estimates for $\Theta = (\rho, \sigma)$. If values had been estimated in a previous analysis, we could have performed a hypothesis test of the observed data using this model as our null hypothesis. Instead, since estimated values are unavailable, the GoF test was used as part of an EDA process to help determine a viable parameter space.

Specifically, the parameter space Θ was estimated using a simulated annealing algorithm (Press et al. 2002: Chapter 10), wherein the GoF test was performed at each different parameterization of the model, and the rank of u_1 was used as the energy function. Again, the rank is indicative of how far the calculated statistic for the observed pattern differs from the mean result for the hypothesized process.

For this example, we selected values of $t_{\min} = 0$ and $t_{\max} = 0.10$, equivalent to an actual distance of 10 m, to represent the distances over which interaction between trees may occur. Additionally, we set the number of simulated patterns within the null distribution to be 499, which yields an estimated value of uncertainty for P values near 0.10 of $\sigma_p = 0.015$. To be clear, only the rank was used here and no hypothesis testing was performed.

The results of the simulated annealing optimization suggest that $\Theta_G = (\hat{\rho} = 42, \sigma = 0.148)$ and $\Theta_K = (27, 0.161)$. The estimated $\hat{\sigma}$ parameters suggest a similar distance of clustering, and the $\hat{\rho}$ parameters are within a similar magnitude.

Discussion

This paper details a statistical test to be used with the G or K point pattern spatial statistics that yields expected type I error rate performance. It is not, however, intended to be a user's guide for evaluating observed point patterns. As such, guidelines about how to choose an appropriate null spatial model or how to select reasonable values for the interaction distance limits ($t_{\min}$ and $t_{\max}$) are outside the current scope. We would like to emphasize, however, that use of CSR as a null model is uninformative in that no spatial processes are truly random, and refer readers to Diggle (2003) for information about how to select something more useful. Our proposed GoF test avoids the problem of simultaneous inference inherent in the simulation envelope approach by converting the results calculated at different distances into a single summary statistic. While the proposed method is certainly not the only solution to this problem, of importance here is that statistical tests based on the GoF approach maintain the expected type I error rate.

We also explain how to estimate the uncertainty in the observed P value as a function of the number of simulations comprising the null distribution. Researchers may find uncertainty in P values somewhat unsettling. Although the GoF test always provides α-level performance, how should we deal with it if the CI around $\hat{p}$ contains our chosen α level and how should we choose the width of this CI? It depends upon the type of inference that is ultimately desired from the analysis. Under strict inference, i.e., when the null hypothesis exists before the data is evaluated and a statistical test is only intended to fail to reject this H_0, a more conservative approach is needed. As such, we suggest that one reject H_0 if the CI is fully below the desired α level. This situation might warrant using a narrower CI based on more simulations and also a tighter type I error rate. Alternatively, under EDA conditions, more latitude can be given. Fewer simulations may be necessary, and it may be justified to increase α levels somewhat, given the sparse nature of data usually collected. Also in this case, we recommended rejecting H_0 if the CI contains or is fully below the specified α level.

Establishment model

Our example, although simplified, demonstrates the major steps to be undertaken during an EDA process, and how the GoF test could be used to facilitate model fitting. A more complete analysis would also test alternate model forms and evaluate the power (or type II error rate) of the hypothesized model. We should clarify that the simulated annealing approach is one of many possible optimization techniques, and under any such method the GoF test can be useful as a cost function. Also, EDA was performed because neither an appropriate spatial model nor parameter estimates were available for the observed data. If both had existed, the

GoF test would have been used to perform a hypothesis test.

Concerning the actual estimated parameters, we found that both the G and K statistics suggest the Poisson cluster model may potentially be an appropriate model for our observed data. However, we also found that parameterization of this model was different for the G and K. These statistics evaluate fundamentally different aspects of a pattern and as such it may not be expected that they would yield identical results. The next step in any such analysis should be an interpretation of the results, including what the fitted spatial model suggests about the ecology of the system. We have learned that a clustered process for the given tree locations is a potentially valid spatial model, but we do not yet have enough information to determine the exact mechanism of clustering. With respect to scale, we have been able to identify a plausible estimate of the range of the clustering, based on the σ model parameter. A next step might examine the species and ages of individual trees to see if clustering is related to establishment or potentially some other environmental condition.

Evaluating scale

By nature, spatial point pattern statistics are multivariate, containing different values at each distance, and it is here that information about scale (i.e., the distances over which a process is acting) is contained. In fact, part of the allure of using the G and K statistics is their purported ability to provide information on the scale of departure of the observed pattern when tested against the hypothesized spatial model. Unfortunately, as described above, because of the lack of statistical significance and being based upon a cumulative value, inference about scale based upon any such departure is suspect.

An improved approach might be built using the derivative of the cumulative result. The u_1 value contains information about the scales that caused the pattern to be rejected. The full range of distances could be partitioned back into range subsets, for multiple GoF tests using different $t_{\min}$ and $t_{\max}$ values to be run. These results could then be evaluated either individually or together using some type of Bonferroni correction. Similar ideas have been suggested by Condit et al. (2000) and Wiegand and Moloney (2004), but we do not know of a rigorous statistical evaluation, including type I and type II error rates, of this approach. As such, we remain cautious about its utility.

Further, there are questions about how to relate the idea of scale to an ecological process that may have caused it, particularly in the absence of discussing some type of model and associated parameters. The concept of scale is understandable with simple spatial models such as hard core inhibition or clustering, i.e., where the radius parameter strictly limits where neighbors can lie. However, it becomes more ambiguous when soft core models such as the Poisson cluster model described above are used. In our view, the best approach toward evaluating scale is only within the context of model parameters.

Software implementation of the GoF test

We believe that ecologists fail to use proper methods because the implementation details provided above were previously unavailable and because common software packages, such as the spatial model in S+ and the splancs and spatstat libraries in R, lack the functions necessary to implement these procedures. Hence, the proposed goodness-of-fit test has been implemented using R software (*available online*).[4] Source code containing functions to test observed patterns via these methods and a user's guide can be found in the Supplement. The current version of software only accommodates square plots and only implements the reduced sample edge correction method. We intend to continue development, and as it is available as open source software, would encourage others to tailor it to their specific needs.

Acknowledgments

We are thankful to Professor Peter Diggle for his review of our proposed methods and discussion of more informative H_0, and to Professor Elizabeth Thompson for her suggestions about the theoretical estimator of the variance associated with the number of simulations. We are also grateful to the WRCCRF staff for allowing us to use this data set.

Literature Cited

Batista, J., and D. Maguire. 1998. Modeling the spatial structure of topical forests. Forest Ecology and Management **110**:293–314.

Casella, G., and R. Berger. 2002. Statistical inference. Duxbury, Pacific Grove, California, USA.

Condit, R., et al. 2000. Spatial patterns in the distribution of tropical tree species. Science **288**:1414–1418.

Cressie, N. A. 1991. Statistics for spatial data. Wiley, New York, New York, USA.

Diggle, P. 2003. Statistical analysis of spatial point patterns. Arnold, London, UK.

Dolezal, J., P. Stastna, T. Hara, and M. Srutek. 2004. Neighbourhood interactions and environmental factors influencing old-pasture succession in the Central Pyrenees. Journal of Vegetation Science **15**:101–108.

Druckenbrod, D., H. Shugart, and I. Davies. 2005. Spatial pattern and process in forest stands within the Virginia piedmont. Journal of Vegetation Science **16**:37–48.

Freeman, E. A., and E. D. Ford. 2002. Effects of data quality on analysis of ecological pattern using the $\hat{K}(d)$ statistical function. Ecology **83**:35–46.

Gotelli, N., and A. Ellison. 2004. A primer of ecological statistics. Sinauer Associates, Massachusetts, USA.

Grassi, G., G. Minotta, G. Tonon, and U. Bagnaresi. 2004. Dynamics of Norway spruce and silver fir natural regeneration in a mixed stand under uneven-aged management. Canadian Journal of Forest Management **34**:141–149.

Harris, M. 2004. The importance of competition processes and canopy gaps in the development of old-growth *Pseudotsuga*/*Tsuga* forests. Thesis. University of Washington, Seattle, Washington, USA.

[4] ⟨http://cran.r-project.org/⟩

Hirayama, K., and M. Sakimoto. 2003. Regeneration of *Cryptomeria japonica* on a sloping topography in a cool-temperate mixed forest in the snowy region of Japan. Canadian Journal of Forest Management **33**:543–551.

Kenkel, N. C. 1988. Pattern of self-thinning in jack pine: testing the random mortality hypothesis. Ecology **69**:1017–1024.

Marriott, F. 1979. Barnard's Monte Carlo tests: how many simulations? Applied Statistics **28**:75–77.

Martens, S., D. Breshears, D. Meyer, and F. Barnes. 1997. Scales of above-ground and below-ground competition in a semi-arid woodland detected from spatial pattern. Journal of Vegetation Science **8**:655–664.

Moeur, M. 1993. Characterizing spatial patterns of trees using stem-mapped data. Forest Science **39**:756–775.

Moeur, M. 1997. Spatial models of competition and gap dynamics in old-growth *Tsuga heterophylla*/*Thuja plicata* forests. Forest Ecology and Management **94**:175–186.

North, M., J. Chen, B. Oakley, B. Song, M. Rudnicki, A. Gray, and J. Innes. 2004. Forest stand structure and pattern of old-growth western hemlock/Douglas-fir and mixed conifer forests. Forest Science **50**:299–311.

Parish, R., J. Antos, and M. Fortin. 1999. Stand development in an old-growth subalpine forest in southern interior British Columbia. Canadian Journal of Forest Research **29**:1347–1356.

Press, W., S. Teukolsky, W. Vetterling, and B. Flannery. 2002. Numerical recipes in C, the art of scientific computing. Second edition. Cambridge University Press, Cambridge, UK.

Rigg, L. 2005. Disturbance processes and spatial patterns of two emergent conifers in New Caledonia. Austral Ecology **30**:363–373.

Salvador-Van Eysenrode, D, J. Bogaert, P. Van Hecke, and I. Impens. 2000. Forest canopy perforation in time and space in Amazonian Ecuador. Acta Oecologica **21**:285–291.

Srutek, M., J. Dolezal, and T. Hara. 2002. Spatial structure and associations in a *Pinus canariensis* population at the treeline, Pico del Teide, Tenerife, Canary Islands. Arctic, Antarctic and Alpine Research **34**:201–210.

Tirado, R., and F. Pugnaire. 2003. Shrub spatial aggregation and consequences for reproductive success. Oecologia **136**: 296–301.

Van Pelt, R., and J. F. Franklin. 2000. Influence of canopy structure on the understory environment in tall, old-growth, conifer forests. Canadian Journal of Forest Research **30**: 1231–1245.

Wiegand, T., F. Jeltsch, I. Hanski, and V. Grimm. 2003. Using pattern-oriented modeling for revealing hidden information: a key for reconciling ecological theory and application. Oikos **100**:209–222.

Wiegand, T., and K. Moloney. 2004. Rings, circles and null-models for point pattern analysis in ecology. Oikos **104**:209–229.

SUPPLEMENT

R source code and user's guide for the implementation of the statistical test described in this paper (*Ecological Archives* E087-120-S1).