ELSEVIER Computational Statistics & Data Analysis 35 (2001) 475–485

COMPUTATIONAL STATISTICS & DATA ANALYSIS

www.elsevier.com/locate/csda

Importance sampling for spatial scan analysis: computing scan statistic p-values for marked point processes ☆

Carey E. Priebe *, Daniel Q. Naiman, Leslie M. Cope

Department of Mathematical Sciences, Whiting School of Engineering, Johns Hopkins University, Baltimore, MD 21218-2682, USA

Received 1 January 2000; received in revised form 1 May 2000

Abstract

Each point in an observed point pattern representing potential target detections (e.g., mines for minefield detection and localization) often is accompanied by a scalar 'mark' representing the detector's level of confidence in that particular detection. Scan analysis for clustering should take this additional mark information into account. We present an importance sampling method for deciding, based on an observed marked point pattern, if a scan statistic S provides significant evidence of increased activity in some localized region of time or space. Our method allows consideration of scan statistics based simultaneously on multiple scan geometries. Our approach yields an unbiased p-value estimate of the form $P[S \geq s_{\text{observed}}] = B^*\rho$, where B^* plays the role of the Bonferroni upper bound and the correction factor ρ measures the conservativeness of this upper bound. The variance of our importance sampling estimate is typically smaller than that of the naive hit-or-miss Monte Carlo technique when the p-value is small. Furthermore, our estimate is often accurate for critical values which are not far enough in the tails of the null distribution to allow for accurate approximations via extreme value theory. In this article, we develop our importance sampling p-value estimator for the case of marked spatial Poisson processes using multiple scan geometries, and illustrate the approach via application to minefield reconnaissance. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Spatial scan analysis; Importance sampling; Marked point pattern; Minefield

☆ This work is partially supported by Office of Naval Research Grant N00014-95-1-0777.

* Corresponding author. Tel.: +1-410-516-7200; fax: +1-410-516-7459.
E-mail address: cep@jhu.edu (C.E. Priebe).

PII: S 0167-9473(00)00017-7

1. Motivation: minefield reconnaissance

Minefield detection and localization is an important problem currently receiving much attention in the engineering and scientific literature (Smith, 1995; Witherspoon et al., 1995). Witherspoon et al. (1995) depict the operational concept for minefield reconnaissance via an unmanned aerial vehicle. Multispectral imagery of an area of interest is processed. First, potential mines are located with a mine detection algorithm. (Holmes et al. (1995) present a thorough discussion of the particular mine detection algorithm used in the sequel.) The detector produces a binary detection map $D(\cdot)$ such that $D(x)=1$ for all points x in the image domain $\mathscr{I} \subset \mathbb{R}^2$ at which a mine or minelike object is detected. Categorizing the candidate detections into 'true targets' (mines) and 'false targets' (mine-like objects, debris, noise, etc.) and considering an operational imperative imposed on the mine detector to find (nearly) all true targets, it can be expected that the number of false targets in the detection map $D(\cdot)$ will be relatively high.

Among the most promising approaches to the minefield detection problem are statistical methods which consider the map $D(\cdot)$ of candidate detections to be a realization of a spatial point process. These methods proceed by analyzing the detection map for clustering or regularity to determine if it represents a minefield point pattern buried in noise, or noise alone (Earp et al., 1995; Muise and Smith, 1995; Cressie and Lawson, 1997; Hayat et al., 1997; Lake et al., 1997; Priebe et al., 1997a).

It is common to consider the detection of clustering in point processes as a formal hypothesis test, with the null hypothesis represented by homogeneity, or complete spatial randomness (Diggle, 1983; Cressie, 1993). For the minefield detection application, in which minefields and false detections are modelled as point processes, the null hypothesis is 'no minefield'. While there are cases for which regularity – patterns in the observed point process – can be used as a key to minefield detection, other applications do not allow for this restriction of the alternative hypothesis of 'minefield present'. It is these latter situations, in which minefield detection becomes a test of homogeneity against a more general alternative of nonhomogeneity, that we address here.

2. Spatial scan analysis

Consider a point process $D : \mathscr{I} \to \{0,1\}$ where the domain $\mathscr{I}$ is a bounded subset of $\mathbb{R}^d$. The goal is to perform a Poisson dispersion test of $\mathrm{H}_0 : \mathscr{P}(\lambda_0)$ on $\mathscr{I}$ for $\lambda_0 > 0$ fixed and known (homogeneity) versus the nonhomogeneity alternative that some subregion has higher intensity. E.g., under H_A the region of interest $\mathscr{I}$ is partitioned into disjoint regions $\mathscr{I}_0$ and $\mathscr{I}_A$ and the point process is $\mathscr{P}(\lambda_0)$ on $\mathscr{I}_0$ and $\mathscr{P}(\lambda_0 + \lambda_A)$ on $\mathscr{I}_A$ with $\lambda_A > 0$. Thus, the map $D(\cdot)$ takes the value zero except at finitely many locations x_i, $1,\dots,N$, at which locations the value taken is unity. The null hypothesis further implies that $N \sim \mathscr{P}(\lambda_0)$ and, conditionally given N, $x_1,\dots,x_N$ are independent and identically distributed uniformly in $\mathscr{I}$. For our application, $\mathscr{I}$ represents the entire area scanned, $\mathscr{I}_0$ is the subregion with no minefield, and $\mathscr{I}_A$

corresponds to the location of the minefield. The intensity of the false alarm process is λ_0, the intensity of the minefield process is λ_A, and $\lambda_0 + \lambda_A$ is the intensity of the compound process (superimposed false alarm and minefield processes) in $\mathscr{I}_A$.

An intuitive approach to testing these hypotheses involves the quadrat counts of Fisher et al. (1922); see Diggle (1983). The generalization to spatial scan statistics is considered in Naus (1965), Cressie (1977,1980), Adler (1984), and Loader (1991), and more recently in Alm (1997), Kulldorff (1997), Chen and Glaz (1997), and Priebe et al. (1997b). A window $W_g(x)$ with geometry (size and shape) g is positioned at each point $x \in \mathscr{I}_g = \{y \in \mathscr{I}\colon W_g(y) \subset \mathscr{I}\}$ and the *locality statistic* $N_{x,g} = \sum_{x' \in W_g(x)} D(x')$ is the number of events observed in the window. The summation is over all $x' \in W_g(x)$; there is (almost surely) a finite number of $x' \in W_g(x)$ such that $D(x') = 1$. The *scan statistic* S is defined to be the maximum over all $x \in \mathscr{I}_g$ of the $N_{x,g}$ and H_0 is rejected for large values of S.

Inference using this scan statistic is difficult; as noted by Cressie (1993), the exact distribution of the statistic for $d \geq 2$ has proved elusive. Much of the available analysis involves extreme value theory, and is therefore applicable only in the tail of the null distribution (for extremely small significance levels). For more moderate significance levels, naive Monte Carlo simulation is an option; however, this approach is computationally intensive, as only a small percentage of generated observations will cause a rejection and thus many repetitions are required in order to reduce the variance of the estimated significance to an acceptable level.

2.1. Multiple scan geometries

An additional complexity arises due to the requirement to consider simultaneously *multiple scan window geometries* $g \in \mathscr{G}$. This consideration allows quantitative inference which takes into account the fact that the size and shape of nonhomogeneities (minefields) are not known a priori and thus the search for a region $\mathscr{I}_A$ of excess intensity must be based simultaneously on multiple window geometries. Loader (1991), Kulldorff (1997), and Naiman and Priebe (2000) have previously considered this important issue.

2.2. Marked point processes

In this article we address an additional consideration; namely, *marked* point processes (see, e.g., Cressie, 1993), in which each detection in the detection map $D(\cdot)$ has associated with it a mark indicating the confidence of the detector in that particular detection. In this case, the map $D(\cdot)$ takes value zero except at finitely many locations x_i, $1, \ldots, N$; at these non-zero locations the value taken is $M(x_i)$. In addition to assuming that the non-zero locations form a homogeneous Poisson process, our null hypothesis now also assumes a simple marking distribution; for instance, that the marks $M(x_i)$ are *iid* uniform random variables ($F_0 = \mathscr{U}[0, 1]$) and are independent of the associated marginal spatial point process. We also assume that, under the alternative, the marks $M(x_i)$ associated with true targets have a distribution F_A which

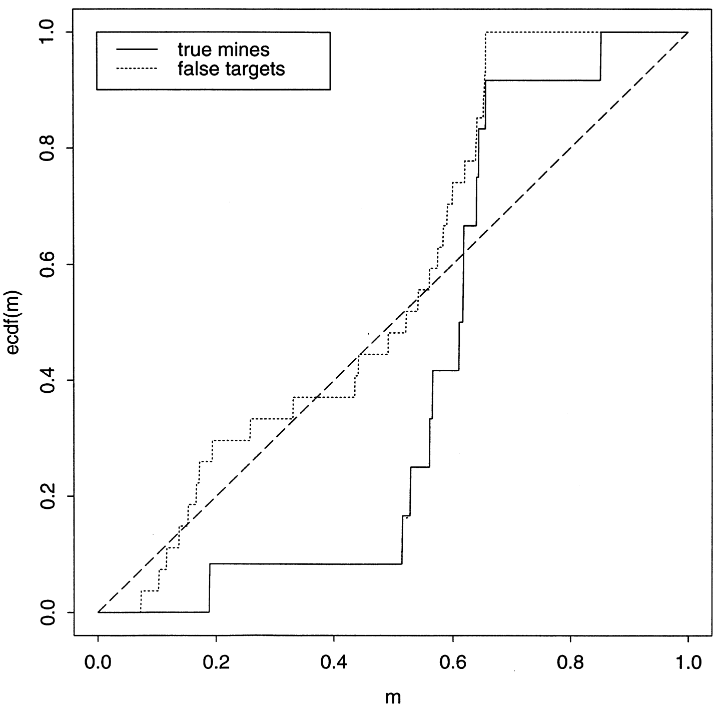

Fig. 1. Empirical cdfs for the class-conditional marks for a minefield data example.

is stochastically larger than F_0. Fig. 1 presents empirical cumulative distribution functions for the class-conditional marks for example minefield data (Witherspoon et al., 1995; Holmes et al., 1995). For our minefield application, these marks correspond to posterior probabilities that the individual detections obtained from the mine detector are in fact mines, as determined by a post-processing classification rule (see, e.g., Priebe et al., 1999). The stochastic ordering assumption is supported, and the uniform assumption for the null marks appears conservative. Incorporation of such marking information will improve the ability to detect nonhomogeneity.

For the marked case with multiple scan geometries the *locality statistic* for $x \in \mathscr{I}_g$ and $g \in \mathscr{G}$ is

$$N_{x,g} = \sum_{x' \in W_g(x)} M(x') - c_g.$$

The *scan statistic* is

$$S = \max_{g \in \mathscr{G}} \max_{x \in \mathscr{I}_{(g)}} N_{x,g}.$$

The geometry-specific constants c_g are determined so that the significance level for the test based on an individual $N_{x,g}$ is the same across window geometries.

3. Importance sampling

Naiman and Priebe (2000) introduce an *importance sampling* approach to estimating the significance level of hypothesis tests based on scan statistics. The general idea of importance sampling (see Fishman, 1996) is to change the distribution to be sampled from, and follow this by making an *adjustment* that accounts for the change to produce an unbiased estimator of the desired probability. Naiman and Priebe's approach to importance sampling in the scan statistic setting involves sampling only data in which rejection of the null hypothesis occurs, and then the subsequent adjustment step involves determining the collection of all locality statistics producing a rejection. There are several benefits of this aproach. First, in a variety of examples, the importance sampling is relatively easy to implement for conveniently structured situations. Second, the importance sampling is more computationally efficient than naive Monte Carlo simulation when the significance levels are sufficiently small, but not so small as to allow for accurate extreme value approximation. Finally, the Naiman and Priebe importance sampling technique is capable of handling multiple scan window geometries $g \in \mathscr{G}$. This article presents the extension of the Naiman and Priebe importance sampling approach to the case of marked point processes with multiple scan window geometries.

Given an observed marked point pattern (obtained through the processing of imagery, for example) and a collection of geometries to be considered, the observed value of the scan statistic, denoted s_{obs}, is calculated. Our importance sampling methodology for estimating the p-value of the observed s_{obs} is designed to improve the naive hit-or-miss Monte Carlo simulation. For the naive simulation, point patterns are generated under the null distribution and scan statistics $S_1,\ldots,S_m$ are observed. The estimate $\tilde{p} = (1/m)\sum_{j=1}^{m} I_{\{S_j \geq s_{\text{obs}}\}}$ can have an unacceptably large variance for computationally reasonable values of the point pattern sample size m when the true p-value is small.

For a scan statistic involving a finite number of locality statistics, our importance sampling approach can be interpreted as providing a correction factor ρ to the conservative *Bonferroni upper bound* $B = \sum_{g\in\mathscr{G}}\sum_{x\in\mathscr{I}_g} P[N_{x,g} \geq s_{\text{obs}}]$ for $P[S \geq s_{\text{obs}}]$. Because the scan statistic under consideration involves a continuum of locality statistics, this interpretation does not make sense; $B = \sum_{g\in\mathscr{G}}\sum_{x\in\mathscr{I}_g} P[N_{x,g} \geq s_{\text{obs}}] = \infty$. However, playing the role of B in our case is $B^* = \sum_{g\in\mathscr{G}} |\mathscr{I}_g| P[N_{x,g} \geq s_{\text{obs}}]$, and we have $P[S \geq s_{\text{obs}}] = B^*\rho$. Our importance sampling requires the ability to efficiently generate sample point patterns from the conditional distribution given *exceedence*, the event $\{S \geq s_{\text{obs}}\}$ that the geometry-normalized sum of the marks in some window equals or exceeds s_{obs}. The procedure provides an estimate $\hat{\rho}$ of the the correction factor ρ which measures the conservativeness of B^*. That is, point patterns conditioned on the event $\{S_j \geq s_{\text{obs}}\}$ are generated for $j = 1,\ldots,m$. For the jth generated conditional point pattern, the set of x's for which $N_{x,g} \geq s_{\text{obs}}$ is determined and the Lebesgue measure of this set, denoted $\hat{\gamma}_j$, is calculated. Then $\hat{\rho} = (1/m)\sum \hat{\gamma}_j^{-1}$.

Our importance sampling estimate $\hat{p} = B^*\hat{\rho}$ is unbiased for the true p-value $P[S \geq s_{\text{obs}}] = B^*\rho$ with a variance that is typically smaller than that of the naive Monte Carlo estimate when the p-value is small. Furthermore, our estimate can be accurate for critical values which are not far enough in the tail of the null distribution to allow for accurate approximations via extreme value theory.

When addressing marked Poisson processes rather than standard unmarked processes, additional complexity arises due to the need to generate *marked* point patterns conditionally on exceedence. This involves four steps. First, the geometry and location of a window W_E in which exceedence is forced to occur are chosen at random; the geometry g_E according to the appropriate weighted distribution and the location x_E uniformly in $\mathscr{I}_{g_E}$. Second, the number of points n for the window W_E is generated conditionally on exceedence in W_E (the event $\{N_{x_E,g_E} \geq s_{\text{obs}}\}$). Then a value s for the sum of the marks in the window is generated conditionally on exceedence in W_E and the event $\{N = n\}$. Finally, marks for the points must be generated from F_0 conditionally on the events $\{N_{x_E,g_E} \geq s_{\text{obs}}\}$, $\{N = n\}$, and $\{\sum_{i=1}^{n} M_i = s\}$. Generating the remainder of the marked point pattern is straightforward, as the locations of points outside W_E are independent of what happens inside W_E, and the marks on the points outside W_E are iid F_0. This algorithm is presented in detail in the appendix.

The trade-off between importance sampling and naive Monte Carlo is one of complexity of sample generation vs. number of samples required to obtain a prescribed precision for the estimate of the p-value. Thus efficient code for the conditional sampling scheme described above is imperative; if the conditional sampling scheme is fast relative to the benefit in terms of precision, then our importance sampling methodology will be preferred. To quantify this comparison, we define the relative efficiency R to be

$$R = (\tilde{t}/\hat{t})(\tilde{\sigma}^2/\hat{\sigma}^2),$$

where $\tilde{t}$ is the elapsed time for the naive estimate and $\hat{t}$ is the elapsed time for the importance estimate, $\tilde{\sigma}^2$ and $\hat{\sigma}^2$ are estimates of the p-value variance for the two procedures, and both algorithms have been run for an equal number of trials. An observed value of R greater than unity favors importance sampling.

The utility of importance sampling for a given application follows from identifying a range of p-values (p^{**}, p^{*}), where p^{**} is the p-value below which importance sampling cannot provide a better estimate than extreme value theory with an acceptable computational load and p^{*} is the p-value above which naive Monte Carlo outperforms importance sampling. If it can be demonstrated that $p^{**} < p^{*}$, and that (p^{**}, p^{*}) encompasses an operationally important range of p-values, then the importance sampling approach should be adopted.

4. Example results

To compare importance sampling vs. naive Monte Carlo for nonhomogeneity detection in marked point processes, and to investigate the utility of using marks, we consider the following scenario motivated by the minefield reconnaissance application (see Fig. 2).

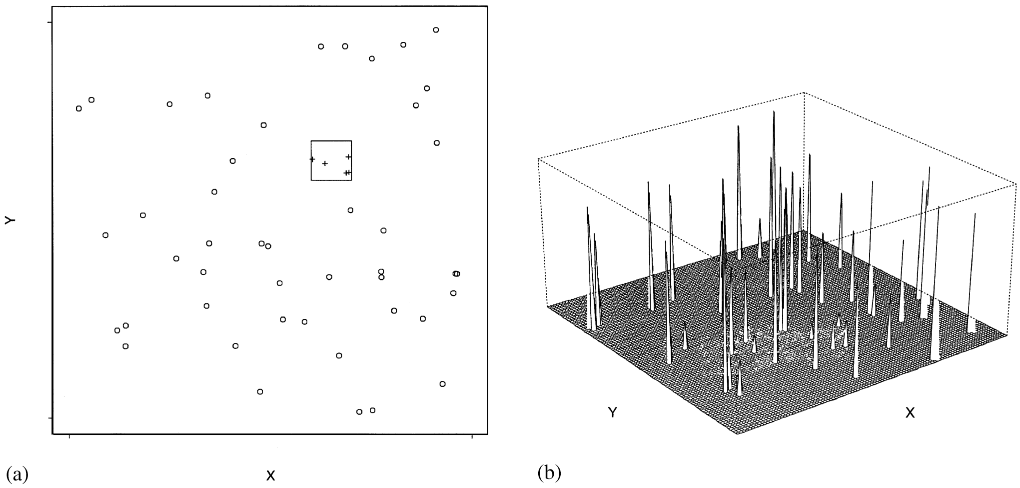

Fig. 2. Marked point pattern for the minefield simulation example. Five "true mines" are located in one scan window, with 50 false alarms expected in the area of interest. Marks (right panel) are uniformly distributed for the false alarms, and stochastically larger for the "true mines".

Let $\mathscr{I}$, the area of interest scanned, be the unit square $[0,1]\times[0,1]$. Let the set of geometries under consideration be $\mathscr{G}=\{g_1,g_2,g_3,g_4\}$ where the g_i are squares with sides of length $l=0.05,0.10,0.15,0.20$ for $i=1,2,3,4$. That is, we suspect that, if it exists, the region of nonhomogeneity $\mathscr{I}_{\mathscr{A}}$ (the minefield) occupies between 0.25% and 4% of $\mathscr{I}$. Assume also that the detection algorithm is calibrated and we expect approximately 50 false alarms in $\mathscr{I}$; that is, $\lambda_0=50$.

For an observed point pattern, our scenario involves consideration of one scan window with $l=0.10$ in which 5 points are observed to occur, and the observed value of the locality statistic for this window determines the observed value of the scan statistic s_{obs}. (These 5 points represent 5 true targets and no false targets, say. That is, we assume the worst case scenario, from a detection point of view, in which no false targets contribute to the value of s_{obs}.) Consider also the analogous scenario, with marks. The marks are iid $\mathscr{U}(0,1)$ under H_0. Our dominating window with $l=0.10$ yields an observed locality statistic of 4. (This represents 5 true targets with marks from F_A averaging 0.8, say.)

Simulation results for this case, based on 1000 Monte Carlo replicates of 1000 trials each, are as follows. For naive Monte Carlo estimation $\tilde{p}=0.0196$ with estimated variance $\tilde{\sigma}^2=1.94\times10^{-5}$, while importance sampling estimation yields $\hat{p}=0.0195$ with estimated variance $\hat{\sigma}^2=1.33\times10^{-5}$. The two estimates agree with one another, and the importance sampling estimate is more precise ($\tilde{\sigma}^2/\hat{\sigma}^2=1.46$). Efficiency comparison of the two estimation algorithms indicates conclusively that, for this case, importance sampling is superior. A sign test of H_0: median$(R)\le 1$ yields a significance value of essentially zero; $R>1$ for 653 of the 1000 trials.

Failure to use the marks would be a grave error for this scenario; the unmarked analysis yields a p-value of $p>0.25$ for this case, compared to $p\approx0.02$ for the

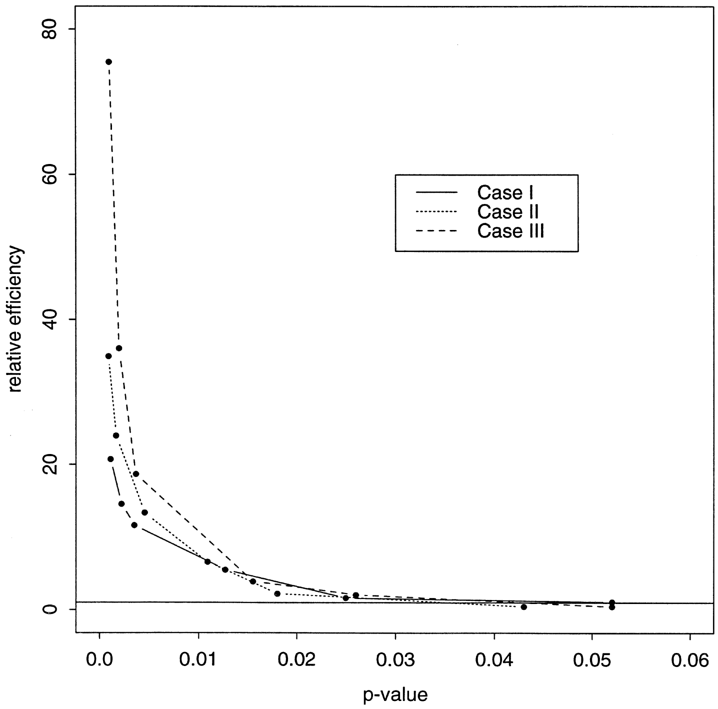

Fig. 3. The relative efficiency of importance sample vs. ñaive Monte Carlo is studied for three cases. Case I considers four geometries simultaneously, with an expected false alarm rate of 10. Case II considers two geometries simultaneously, with an expected false alarm rate of 10. Case III considers two geometries simultaneously, with an expected false alarm rate of 20. We see that, in each case, importance sampling begins to outperform naive Monte Carlo for p-values near 0.05, with the performance improvement becoming quite dramatic as the p-value decreases.

marked analysis. (Naive sampling is more efficient than importance sampling for this unmarked case, as would be expected due to the large p-value.)

Finally, we note that extreme value theory is not a viable option for the scenario under consideration. Even if such a theory can be developed for the case of marked point processes with multiple scan geometries, the observed $p \approx 0.02$ is not likely to be far enough in the tail of the null distribution to allow for accurate estimation via extreme value theory.

Thus, for our example scenario we have demonstrated the utility of importance sampling: $0.02 \in (p^{**}, p^{*})$. Fig. 3 presents the results of further simulation studies, for a range of p-values and multiple choices of geometries $\mathscr{G}$ and background intensities λ_0. For Case I we set $\lambda_0 = 10$ and, as in the previous example, four geometries are considered; $\mathscr{G} = \{g_1, g_2, g_3, g_4\}$ where the g_i are squares with sides of length $l = 0.05, 0.10, 0.15, 0.20$ for $i = 1, 2, 3, 4$. Case II considers just two geometries, $\mathscr{G} = \{g_1, g_4\}$, with $\lambda_0 = 10$ as in Case I. Case III considers the same two geometries as Case II, with an increased background intensity of $\lambda_0 = 20$. The simulation results suggest that, for each of these three cases, importance sampling begins to outperform naive Monte Carlo for p-values near 0.05, and that the performance improvement is quite significant for p-values less than 0.02 and dramatic for p-values less than 0.01. Analysis indicating the extent to which these results generalize is of obvious interest and importance and is the subject of ongoing research.

5. Summary and conclusions

In this paper marked spatial Poisson processes and spatial scan analysis involving multiple scan geometries is considered. An importance sampling algorithm for estimating p-values for a test of homogeneity is presented. For a scenario motivated by a minefield detection example it is demonstrated that (1) the importance sampling estimate is superior to naive Monte Carlo estimation in terms of relative efficiency, and (2) it is operationally imperative to use the marking information. Furthermore, there does not currently exist a viable extreme value theory for estimation in the scenario considered. The conclusion is that the importance sampling approach presented is the methodology of choice for p-value estimation in the scenario considered.

In addition to the ongoing quest for analytic results regarding the range (p^{**}, p^{*}) for various general scenarios, extensions of this importance sampling approach are currently being investigated. The binomial process can be treated as above. Inhomogeneous Poisson processes and Cox processes are of significant interest in the minefield detection application and can be addressed in much the same way. Spatial point processes exhibiting dependency (Poisson cluster processes, inhibition processes, Markovian processes) are also of operational interest. Finally, more complex marking distributions are being considered.

Appendix Importance sampling algorithm

Although it is the final result of a scan using all $g \in \mathscr{G}$, s_{obs} will be associated with one specific, *dominant* geometry g_s. In order to estimate p-values, we need an equivalent threshold value t_g for each g.

Fixing $x \in \mathscr{I}_{g_s}$, define $\alpha := P[N_{x,g_s} \geq s_{\text{obs}}]$. For each $g \in \mathscr{G}$, fix $y \in \mathscr{I}_g$ and let t_g be such that $P[N_{y,g_s} \geq t_g] = \alpha$. In the marked case, the distribution is continuous and t_g is unique.

Step 1. *Randomly select geometry and location of window W_E.*

Select $g_E \in \mathscr{G}$ according to the probabilities

$$p_g = \frac{|\mathscr{I}_g| P[N_{x,g} \geq t_g]}{\sum_{g \in \mathscr{G}} |\mathscr{I}_g| P[N_{x,g} \geq t_g]} = \frac{|\mathscr{I}_g|}{\sum_{g \in \mathscr{G}} |\mathscr{I}_g|}.$$

$P[N_{x,g} \geq t_g] = \sum_{n \geq 0} P[W_g(x) \text{ contains } n \text{ points}] P[\sum_{i=i}^{n} U_i \geq_g]$, where $U_i \sim \mathscr{U}[0,1]$.

Choose a location uniformly in $\mathscr{I}_{g_E}$.

Step 2. *Randomly generate the number of points n in W_E, conditioning on exceedence.*

Select n according to the probabilities

$$q_\eta = P[n = \eta] = \frac{P[W_E \text{ contains } \eta \text{ points}] P[\sum_{i=i}^{\eta} U_i \geq t_g]}{\sum_{\eta=1}^{\infty} P[W_E \text{ contains } \eta \text{ points}] P[\sum_{i=i}^{\eta} U_i \geq t_g]}.$$

Step 3. *Generate the sum of marks s, conditional on exceedence, and on the value of n.*

Choose s according to the distribution

$$F_n(s) = P\left[\sum_{i=i}^{n} U_i \le s \,\middle|\, \sum_{i=i}^{n} U_i \ge t_g\right].$$

Step 4. *Generate the individual mark values conditional on the values of s and n.*

Randomly choose a point $x = (x_1 \dots x_n)$ in the simplex

$$\mathscr{S} = \left\{ y\colon\ \sum_{i=1}^{n} y_i = s, y_i \ge 0 \right\}.$$

Place the marked points uniformly in W_E.

Step 5. *Place points outside of* W_E.

The location of points in W_E is independent of the location of points in $\mathscr{I}/W_E$, where one has a Poisson spatial process of intensity λ. Generate iid $\mathscr{U}[0,1]$ marks for these points.

Step 6. *Estimate* γ, *the Lebesgue measure of* $\bigcup_{g\in\mathscr{G}}\{x \in \mathscr{I}_g\colon N_{x,g} \ge t_g\}$.

Methods will be specific to the set of geometries under consideration.

The above steps are performed m times. Define $\hat{\rho}:=(1/m)\sum_{i=1}^{m} \hat{\gamma}_i^{-1}$.

The estimate of the p-value is

$$\hat{\rho}B^* = \hat{\rho}\sum_{g\in\mathscr{G}} |\mathscr{I}_g| P[N_{x,g} \ge t_g] = \hat{\rho}\alpha \sum_{g\in\mathscr{G}} |\mathscr{I}_g|.$$

References

Adler, R.J., 1984. The supremum of a particular Gaussian field. Ann. Probab. 12, 436–444.

Alm, S.E., 1997. On the distributions of scan statistics of a two-dimensional poisson process. Adv. Appl. Probab. 29, 1–18.

Chen, J., Glaz, J., 1997. Two-dimensional discrete scan statistics. Statist. Probab. Lett. 31, 59–68.

Cressie, N.A.C., 1977. On some properties of the scan statistic on the circle and the line. J. Appl. Probab. 14, 272–283.

Cressie, N.A.C., 1980. The asymptotic distribution of the scan statistic under uniformity. Ann. Probab. 8, 828–840.

Cressie, N.A.C., 1993. Statistics for Spatial Data. Wiley, New York.

Cressie, N.A.C., Lawson, A.B., 1997. Models and inference for clustering of locations of mines and minelike objects. Proceedings of the SPIE Vol. 3079: Detection Technologies for Mines and Minelike Targets II, Orlando, FL, April 1997, pp. 519–530.

Diggle, P.J., 1983. Statistical Analysis of Spatial Point Patterns. Academic Press, New York.

Earp, S.L., Elkins, T.J., Conrath, B.C., 1995. Detection of random minefields in clutter. Proceedings of the SPIE Vol. 2496: Detection Technologies for Mines and Minelike Targets, Orlando, FL, April 1995, pp. 543–555.

Fisher, R.A., Thornton, H.G., Mackenzie, W.A., 1922. The accuracy of the plating method of estimating the density of bacterial populations, with particular reference to the use of Thornton's agar medium with soil samples. Ann. Appl. Biol. 9, 325–359.

Fishman, G.S., 1996. Monte Carlo. Springer, New York.

Hayat, M.M., Gubner, J.A., Chang, W.-B., 1997. Detection of signals modeled by interaction point processes using a poisson approximation. Proceedings of the SPIE Vol. 3079: Detection Technologies for Mines and Minelike Targets II, Orlando, FL, April 1997, pp. 531–537.

Holmes, Q.A., Schwartz, C.R, Seldin, J.H., Wright, J.A., Wieter, L.J., 1995. Adaptive multispectral CFAR detection of land mines. Proceedings of the SPIE Vol. 2496: Detection Technologies for Mines and Minelike Targets, Orlando, FL, April 1995, pp. 421–432.

Kulldorff, M., 1997. A spatial scan statistic. Comm. Statist.: Theory Methods 26, 1481–1496.

Lake, D.E., Sadler, B.M., Casey, S.D., 1997. Detecting regularity in mine fields using collinearity and a modified Euclidean algorithm, Proceedings of the SPIE Vol. 3079: Detection Technologies for Mines and Minelike Targets II, Orlando, FL, April 1997, pp. 500–507.

Loader, C.R., 1991. Large-deviation approximations to the distribution of scan statistics. Adv. Appl. Probab. 23, 751–771.

Muise, R.R., Smith, C.M., 1995. A linear density algorithm for patterned minefield detection, Proceedings of the SPIE Vol. 2496: Detection Technologies for Mines and Minelike Targets, Orlando, FL, April 1995, pp. 586–593.

Naiman, D.Q., Priebe, C.E., 2000. Computing scan statistic p-values using importance sampling. J. Comput. and Graphical Statist., to appear.

Naus, J.I., 1965. Clustering of random points in two dimensions. Biometrika 52, 263–267.

Priebe, C.E., Pang, J.S., Olson, T., 1999. Optimizing sensor fusion for classification performance. Proceedings of the CISST '99 International Conference, Las Vegas, Nevada, June 1999, pp. 397–403.

Priebe, C.E., Olson, T.E., Healy, D.M., 1997a. Exploiting stochastic partitions for minefield detection, Proceedings of the SPIE Vol. 3079: Detection Technologies for Mines and Minelike Targets II, Orlando, FL, April 1997, pp. 508–518.

Priebe, C.E., Olson, T., Healy, D.M., 1997b. A spatial scan statistic for stochastic scan partitions. J. Amer. Statist. Assoc. 92, 1476–1484.

Smith, D.L., 1995. Detection technologies for mines and minelike targets. Proceedings of the SPIE Vol. 2496: Detection Technologies for Mines and Minelike Targets, Orlando, FL, April 1995, pp. 404–408.

Witherspoon, N.H., Holloway, J.H., Davis, K.S., Miller, R.W., Dubey, A.C., 1995. The coastal battlefield reconnaissance and analysis (COBRA) program for minefield detection. Proceedings of the SPIE vol. 2496: Detection Technologies for Mines and Minelike Targets, Orlando, FL, April 1995, pp. 500–508.