

An early warning system for space-time cluster detection

RENATO M. ASSUNÇÃO¹, ANDRÉA IABRUDI TAVARES¹, MARTIN KULLDORFF²

¹Laboratório de Estatística Espacial

Universidade Federal de Minas Gerais

assuncao@est.ufmg.br, iabrudi@dcc.ufmg.br

²Biostatistician Department of Ambulatory Care and Prevention

Harvard Medical School and Harvard Pilgrim Health Care

martin.kulldorff@hms.harvard.edu

Abstract. A new topic of great relevance and concern has been the design of efficient early warning systems to detect as soon as possible the emergence of spatial clusters. In particular, many applications involving spatial events recorded as they occur sequentially in time require this kind of analysis, such as fire spots in forest areas as in the Amazon, crimes occurring in urban centers, locations of new disease cases to prevent epidemics, etc. We propose a statistical method to test for the presence of space-time clusters in point processes data, when the goal is to identify and evaluate the statistical significance of localized clusters. It is based on scanning the three-dimensional space with a score test statistic under the null hypothesis that the point process is an inhomogeneous Poisson point process with space and time separable first order intensity. We discuss an algorithm to carry out the test and we illustrate our method with space-time crime data from Belo Horizonte, a large Brazilian city.

1 Introduction

Suppose that data are available consisting of the locations and reference times of events occurring within a specified geographical region and time period. It is common to test whether there is space-time clustering of events, after adjusting for purely spatial and purely temporal clustering. That is, it is of interest to test whether cases which are close in space are also relatively close in time, and vice versa. If so, the data exhibit *space-time clustering*, or in epidemiological terminology, space-time interaction.

The most popular statistical technique for testing space-time interaction with point process data was proposed by Knox (1964). Specifying a spatial and a temporal critical distance, it is possible to indicate when a pair of events is close in space or close in time. The test is based on the number X of pair of events which are simultaneously close in space and in time. A large number X would be an indication that cases which are close in space tend also to be close in time leading to space-time interaction.

Knox test and the later developments by Mantel (1967), Diggle *et al* (1995), Baker (1996), Jacquez (1996), and Kulldorff and Hjalmars (1999) are global tests in that they test for space-time clustering throughout the data, without identifying specific clusters. That is, the tests do not aim at detecting and localizing clusters. This is appropriate when the test is aimed at for example finding evidence of whether a disease is infectious or not. When spatially localized episodic or epidemic outbreak occurs, the identification of clusters is important since the space-time interaction will appear in the form of raised incidence on localized regions

in the space-time volume under study.

Based on the score function, we derive and present a new space-time cluster detection scan statistic for space-time point processes in section 2. In Section 3, we discuss the computer issues involved on the methodology implementation. In Section 4, we apply the methodology to three crime data sets and conclude in Section 5.

2 Identifying Space-Time Clusters

Assume that we observe random point events generated by a Poisson point process in a space-time region $\mathcal{A} = A \times [0, \tau]$, where A is a bi-dimensional polygon. Usually, there is substantial spatial and temporal heterogeneity and this is modelled by the space-time intensity function denoted by $\lambda(x, y, t)$.

Given the observed events, the Poisson log-likelihood is given by

$$l = \sum_{i=1}^n \log \lambda(x_i, y_i, t_i) - \int_{\mathcal{A}} \lambda(x, y, t) dx dy dt$$

The null hypothesis of no space-time interaction implies that the intensity function is a product of two functions, one depending only on the spatial location (x, y) and another depending only on t :

$$H_0 : \lambda(x, y, t) = \lambda_S(x, y) \lambda_T(t)$$

Let $C = C_S \times C_T$ be a fixed and arbitrary space-time cylinder with C_S being a convex region in A and C_T a time in-

terval. Consider a local alternative $H_{C,\epsilon}$ to H_0 given by

$$H_{C,\epsilon} : \lambda(x, y, t) = \lambda_S(x, y) \lambda_T(t) (1 + \epsilon I_C(x, y, t))$$

where $\epsilon > 0$ and I_C is the indicator function that $(x, y, t) \in C$. Therefore, this alternative considers a situation where locally, at the cylinder C , the point process deviates from the space-time separability hypothesis by the larger than expected events density under the null hypothesis (in case $\epsilon > 0$).

The score statistics is given by

$$\frac{\partial l}{\partial \epsilon} \Big|_{\epsilon=0} = N(C) - \int_{C_S} \lambda_S(x, y) dx dy \times \int_{C_T} \lambda_T(t) dt \quad (1)$$

where $N(C)$ is the number of events within C . It can be shown that, under the null hypothesis, (1) becomes

$$\frac{\partial l}{\partial \epsilon} \Big|_{\epsilon=0} = N(C) - \frac{E[N(C_S \times [0, T])] E[N(A \times C_T)]}{E[N(A \times [0, T])]}$$

After some algebraic manipulation, we found that the standardized test statistic is given by

$$U_C = \frac{N(C) - N(C_S \times [0, T]) N(A \times C_T) / N(A \times [0, T])}{\sqrt{N(C_S \times [0, T]) N(A \times C_T) / N(A \times [0, T])}} \quad (2)$$

which is the locally most powerful test in the sense that it maximizes the derivative of the power function at $\epsilon = 0$ (Cox and Hinkley, 1974, page 113).

Usually we have no prior knowledge of space-time clusters location and then the test developed can not be applied since we have no cluster candidate C to use. Hence, our proposed test is based on the scan statistic

$$U = \sup_C \{U_C\} \quad (3)$$

which searches over all possible cylinders C . In practice, the scanning in (3) is undertaken over a smaller class of cylinders for several reasons, not described here for lack of space.

The sampling distribution of U defined in (3) is intractable. As a consequence, its null hypothesis distribution is obtained by a Monte Carlo procedure conditionally on the realizations of the process spatial and temporal components. Under the null hypothesis, the sampling distribution of U is the distribution induced by random permutation of the times $t_i, i = 1, \dots, n$ keeping fixed the spatial locations $(x_i, y_i), i = 1, \dots, n$.

The observed value u_1 of U is ranked amongst values u_2, \dots, u_m generated by recomputing the U statistic after $m - 1$ independent random permutations of the times $t_i, i = 1, \dots, n$. If u_1 ranks k -th largest, the one-sided exact attained significance level is k/m .

This Monte Carlo method is computer intensive and naive algorithms should not be used for large datasets. Algorithmic considerations are discussed in Section 3.

As in Kulldorff (1997), our Monte Carlo test identifies also secondary clusters besides the cylinder C_1 which maximizes the score statistic U defined in (3). That is, it identifies clusters C_2, C_3, \dots non-overlapping with C_1 and with score statistics $U_{C_1} \geq U_{C_2} \geq U_{C_3} \geq \dots$ significantly larger than the $(1-\alpha)$ quantile threshold based on the Monte Carlo reference distribution of $U = \max_C U_C$ under the null hypothesis. We need to consider non-overlapping cylinders because other cylinders, almost coinciding with C_1 , will have a likelihood and score statistic close to that of C_1 . Hence, the second most significant cylinder should not be defined based only on the second ranked cylinder but rather in the second ranked cylinder among those that do not intersect C_1 . It should be noted that the test results for the secondary clusters C_2, C_3, \dots are conservative (Kulldorff, 1997).

3 Computer Implementation Issues

In the scan procedure, we need to consider only the minimum enveloping cylinder of a given subset \mathcal{E} of events. Because $(N(C) - \mu) / \sqrt{\mu}$ is a decreasing function of μ , it is maximized at the minimum value of μ for fixed $N(C)$. Since $C_1 = C_{S1} \times C_{T1} \subset C_2 = C_{S2} \times C_{T2}$ implies that $N(C_{S1} \times [0, T]) N(A \times C_{T1}) \leq N(C_{S2} \times [0, T]) N(A \times C_{T2})$, we have $U_{C1} \leq U_{C2}$. Therefore, it suffices to scan all distinct subsets \mathcal{E} of events and their associated enveloping cylinders.

We additionally restrict the spatial cylinders C_S to be circles centered in an observed event e . Hence, the space-time cylinder $C(\mathcal{E}, e) = (x_e, y_e, r, t_m, t_M)$, centered at $e \in \mathcal{E}$, with space cylinder $C_S(\mathcal{E}, e) = (x_e, y_e, r)$ and time cylinder $C_T(\mathcal{E}, e) = (t_m, t_M)$, is defined by:

$$\begin{aligned} t_m &= \min_{f \in \mathcal{E}} t_f \\ t_M &= \max_{f \in \mathcal{E}} t_f \\ r &= \max_{f \in \mathcal{E}} (d(e, f) = \sqrt{(x_e - x_f)^2 + (y_e - y_f)^2}) \end{aligned}$$

As relevant cylinders have at least s_m events, a naive approach to the scan algorithm is to generate all subsets with size greater than or equal to s_m . This naive approach has $O(n^5)$ order which can be reduced to $O(n^4)$ if we preprocess time and space cylinders, as we explain next.

Space-time cylinders generation is performed by precomputing space and time cylinders. This induces two events orderings that allows cylinders' size evaluation to be performed in constant time ($O(1)$). We iterate through events, get time and space cylinders in the pre-computed orderings and define space-time cylinders by their intersection. We denote by C^* the space-time cylinder whose test statistic value U_{C^*} is maximum. We assign $U_{C^*} = -\infty$ in the beginning of the algorithm. If we let the keyword *loop*-

po be the iteration over a set in a previously specified order, the high-level description of our algorithm is:

Scan Procedure

Input: event set with associated temporal and spatial ranks

Output: C^* and U_{C^*}

SL1. for each event e do

SL2. loop-po over C_T that contains e do

SL3. loop-po over C_S centered in e do

SS1. $C \leftarrow C_S \times C_T$

SS2. compute $N(C), N(C_S), N(C_T)$

SS3. if $N(C) > s_m$ and $U_C > U_{C^*}$, then $C^* \leftarrow C$

return C^* and U_{C^*}

This algorithm has complexity $O(n^4)$, as we show in the Appendix.

3.1 Monte Carlo Procedure

The null distribution of U is obtained conditioning on the observed times and positions of the events. We permute the events' time indexes and apply the scan procedure to obtain the space-time cluster C with largest U_C value. This is repeated independently a large number p of times. Previously computed temporal and spatial ranks are reused. The algorithm for this Monte Carlo procedure is:

Monte Carlo Procedure

Input: event set, temporal and spatial ranks, number p

Output: p -dim array l with largest U_C value in each permutation.

ML1. for $i = 1$ to p do

MS1. generate a permutation of events time index

MS2. update temporal ranks

MS3. call **ScanProcedure**

MS4. $l[p] \leftarrow$ value of U statistic

return l

Considering that generating a random permutation takes $O(n)$ operations, the computational complexity of the procedure is $O(pn^4)$ since loop **ML1** complexity is dominated by step **MS3**.

3.2 Non-overlapping scan algorithm

A first approach to generate non-overlapping best space-time cylinders is to visit each cylinder just one time and to store a list of candidates. As this list may contain up to $O(n^4)$ cylinders, efficient data structures are essential to maintain low computational complexity. Here, we took a different approach since we expect that the number b of desired non-overlapping cylinders is relatively low with respect of the number of relevant cylinders. Instead of storing candidate cylinders, we perform the scan procedure b

times, maintaining a list of already identified best space-time cylinders. This approach reevaluates U_C statistics for many cylinders C but does not require complex data structures. Let G be the list of non-overlapping cylinders. Initially, it is empty and, at each instance of loop **NL1**, a new space-time cylinder is included in G or the procedure terminates if there is no such cylinder.

Non-overlapping Scan Procedure

Input: event set, temporal and spatial ranks and b

Output: $C_1^*, C_2^*, \dots, C_b^*$ and their test statistic U_C values
 $G \leftarrow \emptyset$

NL1. do b times

 found \leftarrow FALSE

 for each event e do

 for each valid C_T that contains e do

 for each valid C_S centered in e do

$C \leftarrow C_S \times C_T$

NS2. if C do not overlap any cylinder in G then

 found \leftarrow TRUE

 compute $N(C), N(C_S), N(C_T)$

 if $N(C) > s_m$ and $U_C > U_{C^*}$, then $C^* \leftarrow C$

 if (found)

 then include C^* in G

 else return G

return G

The complexity of this procedure depends on the number b of desired cylinders, both because of loop **NL1** and step **NS2**. We perform the basic scan b times and, in each one of those, all previously generated cylinders should be inspected in step **NS2**. The final complexity is $O(b^2n^4)$.

4 Space-time clusters in crime data

For illustration, we use the crime incidence data from a large Brazilian city, Belo Horizonte, during 1995-2001 collected by the Polícia Militar de Minas Gerais based on their police records of crime events. Each crime event was georeferenced by the coordinates of its occurrence place (in meters) and occurrence day.

We deal with robberies of three types of stores: drug-stores, bakeries, and lottery houses. Most of these robberies were gun armed robberies. We also consider homicides, a different kind of crime to contrast to the store robberies. We have data from 1995 to 2000 for store robberies and from 1995 to 2001 for homicides. The total number of events were 582, 765, and 2216 for the lottery houses, drugstores, and bakeries robberies, respectively, and 1356 for homicides.

We ran the usual Knox test with two spatial and temporal thresholds: 2 and 3 kilometers and 20 and 30 days. We used 999 Monte Carlo simulations in the tests. Table 1 presents the results. While homicide shows no evidence of




Figure 1: Maps of Belo Horizonte with four types of crime. The upper row shows the 765 drugstore robberies (left) and the 2216 bakery robberies (right). The bottom row shows the 582 lottery house robberies (left) and the homicides (right). The first three range from 1998 to 2000 while homicides data range from 1998 to 2001.

space-time clustering, bakery and drugstore robbery have small p -values while the evidence for lottery robbery depends on the threshold used and even then it is only borderline significant.

To run our scan procedure in the same dataset, we used a minimum of 5 events in each cylinder and we limit the cylinders to contain at most 15% of the observed events and to not cover more than 15% of either the area A or the total time interval $[0, \tau]$. We used 999 Monte Carlo simulations to generate the null hypothesis distribution.

We found C_1^* as a significant (at 0.05 level) space-time cluster in all four crimes, with bakery robberies presenting also C_2^* as a significant cluster (see Table 1). The number of events in the most significant cluster was 5, 7, 6, and 5 events for bakery, drugstore, lottery robberies, and homicide, respectively. The second significant cluster of bakery robberies had 5 events. Although the homicide space-time cluster presented borderline significance, we can see that the scan test identified clusters in homicide and lottery robberies, whereas Knox test did not. This suggests that our method could be more sensitive to the presence of localized clusters than Knox test.

Figure 2 presents the maps from Figure 1 zoomed to show the significant space-time clusters we detected with our scan procedure. The cluster covering the largest area was drugstore robbery with 1.52 kilometers of radius while the other clusters ranged from 370 to 760 meters. Hence, the events within the clusters are tightly clustered in space.

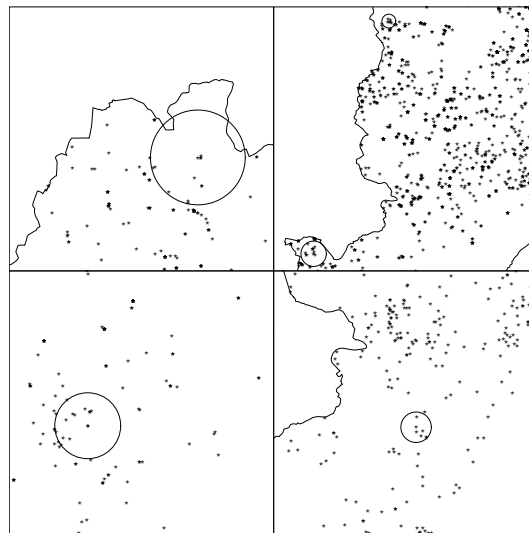


Figure 2: Zoomed maps of Belo Horizonte with the significant space-time clusters C_j^* by type of crime, in the same order as Figure 1. Bakery robbery is the only crime with two significant space-time clusters, the others having only one significant cluster (at 0.05 value).

One store was robbed twice in each of the lottery and drugstore clusters. A more extreme pattern was found in the second bakery cluster with one store robbed three times. There were no obvious spatial pattern connecting clusters from different crimes.

Concerning time, the shortest bursts of spatially localized violence was that associated with the two clusters of bakery robberies. They first and second clusters C_1^* and C_2^* lasted 8 and 17 days starting on February, 28 2000 and March 29, 2000 respectively. Drugstore and lottery robberies had longer clusters lasting 68 and 81 days starting on April 03, 1997 and May 23, 1995, respectively. The homicide cluster was detected on February 03, 2000, lasting 58 days.

The significant clusters of bakery robberies showed extreme patterns. Cluster C_1^* lasted only 8 days and, although occurring in different parts of the city, the second cluster started only 3 weeks after the first one had disappeared. This lasted only 17 days and contained five events related with 3 different stores, one of them being robbed three times during this period and five times during the total study period. The time lags between the five successive events in this second space-time cluster were 5, 2, 6, and 4 days. To improve visualization in Figure 2, the repeated robbed store events coordinates were jugged by random normal random variables with mean zero and standard deviation 60.

Crime	2 km 20 days	3 km 30 days	C_1^*	C_2^*
Bakery robb.	0.01	0.01	0.030	0.032
Drugstore robb.	0.01	0.01	0.012	0.154
Lottery robb.	0.05	0.22	0.028	0.220
Homicide	0.10	0.11	0.048	0.344

Table 1: Table with the p-values of Knox and scan tests. The results are separated according to either the thresholds used in the test (Knox) or the first (C_1^*) and second (C_2^*) most significant cylinders (scan test). The null hypothesis distribution was determined by 999 Monte Carlo permutations of the observed times t_i .

5 Conclusion

Recently, there has been interest on space-time surveillance systems for the early detection of disease outbreaks, but very few studies have provided solutions to this problem. Rogerson (2001) proposed to use a cumulative sum approach for space-time point processes data, each event being scored according to a local Knox statistic. Theophilides (2003) uses multiple Knox test (1964) in multiple local area, so see whether one or more exhibit space-time clustering within that local area. Their method is very different from the one proposed here. Kulldorff et al. (2002,2003) have developed a space-time permutation scan statistic for the early detection of disease outbreaks, which is currently in use by the New York City Department of Health for syndromic surveillance. They use a Poisson based likelihood ratio test statistic rather than the score test proposed in this paper. All of these methods are prospective in nature, in that they are looking for recent outbreaks, as opposed to our retrospective method, which aims at detecting space-time clusters at any location and time.

In conclusion, our method has many desirable features: it does not require population data; it identifies the space-time clusters; it does not require time and distance critical thresholds as Knox test does, it adjusts for purely spatial and purely temporal clustering, and it provides statistical inference for each individual cluster detected. We think it will be of great use in many practical applications.

Acknowledgements

This work was partially supported by project SAUDAVEL, joint call MCT/SEPIN - FINEP - CNPq 01/2002. The dataset used in the example was kindly provided by the Polícia Militar de Minas Gerais.

Appendix: Algorithm Complexity

The algorithm has two steps. The first one finds the relevant time and space cylinders and this takes $O(n \log n)$ and $O(n^2 \log n)$ complexity, respectively. The second step finds the space-time cylinders as intersections of space and time cylinders and this takes $O(n^4)$ running time. Therefore, a single execution of the scan procedure takes

$$O(n \log n) + O(n^2 \log n) + O(n^4) = O(n^4)$$

running time. It takes $O(pn^4)$ to obtain the null hypothesis distribution with p independent permutations of time indexes. In order to obtain the b non-intersecting most significant cylinders, we need $O(b^2 n^4)$ running times. We describe these calculations with more detail.

Temporal and spatial ranks input

Let $t_{(1)} < t_{(2)} < \dots < t_{(n)}$ be the ranked values of the observed times and m_t be the maximum length of a time cylinder. The set of time cylinders \mathcal{C}_T induced by the observed events is

$$\mathcal{C}_T = \{(t_{(i)}, t_{(j)}) \mid i < j \text{ and } t_{(j)} - t_{(i)} \leq m_t\}.$$

We derive the time cylinders from the time ranked values and these are obtained in $O(n \log n)$ running time if we use optimal sorting algorithm, such as *quicksort*.

The set of relevant space cylinders \mathcal{C}_S is the union of the space cylinders centered at events e . Let (x_e, y_e) be the center of a space cylinder and $d_{(2)}^e < d_{(3)}^e < \dots < d_{(n)}^e$ be the ranked distances of the other observed events to event e . Note that $d_{(1)}^e = 0$ since it is the distance between e and e . If m_d is the maximum radius of a space cylinder, the set of space cylinders \mathcal{C}_S is

$$\mathcal{C}_S = \left\{ (x_e, y_e, d_{(s)}^e) \mid e, s \in \{1, \dots, n\}, s \geq 2, \text{ and } d_{(s)}^e \leq m_d \right\}.$$

The distance ranked values are sufficient to define space cylinders. They can be evaluated by n sorting algorithms. Therefore, we need $O(n^2 \log n)$ running time to generate the spatial cylinders.

Scan procedure complexity

Given an event e , loops **SL2** and **SL3** should guarantee the generation of all time cylinders containing e and all space cylinders centered at e . A time cylinder is generated by defining first an allowable initial time $t_m = t_{(i)}$ ($t_m \leq t_e$ and $t_e - t_m \leq m_t$), and then, a valid final time $t_M = t_{(f)}$ ($t_e \leq t_M$ and $t_M - t_m \leq m_t$). For this, we iterate

through events in their temporal ordering determining i and f in $O(\log n)$ time. Loop **SL2** is implemented by letting m vary from i till $r_t(e)$ where $r_t(e)$ is event e temporal rank. For a fixed m , valid time cylinders $C_T = (t_{(m)}, t_{(M)})$ are those that contains e (M varies from $r_t(e)$ to f) and respects time threshold m_t . There are at most $\sum_{i=1}^n (n - i) < n^2$ such cylinders (depending on the m_t value), thus **SL2** is performed $O(n^2)$ times. The size of a time cylinder $C_T = (t_{(m)}, t_{(M)})$ in **SS2** is $M - m + 1$ events and can be evaluated in constant time.

For a fixed time cylinder C_T , space cylinders generated in loop **SL3** should contain at least the m -th and M -th time events in order to minimize the re-generation of space-time cylinders. Let e_m and e_M be those events. The space cylinder minimum radius is $r = \min\{d(e, e_m), d(e, e_M)\}$. Let l be the distance rank of this radius, which can be determined in $O(\log n)$. Spatial cylinders in loop **SL3** are of the form $C_S = \{(x_e, y_e, d_{(s)}^e) \mid l \leq s < g\}$, where $d_{(g)}^e > m_d$ or $g = n + 1$. Loop **SL3** is performed $O(n)$ times. The evaluation of $N(C_S)$ in **SS2** is exactly s , and takes constant time.

The space-time cylinders are the intersection of events in a spatial and a time cylinder. For a fixed C_T in loop **SL2**, spatial cylinders are considered in increasing radius in loop **SL3**. Hence, a new spatial cylinder must include all previous events and at least one additional event. Initially, a sequential scan is used to determine the intersection of C_T and $C_S = (x_e, y_e, d_{(l)}^e)$, by comparing the times t_e of events $e \in C_S$ to the time interval of C_T . For each new C_S , a new event is included and the comparison is repeated. It results that the computation of all C in step **SS3** for a loop **SL3** is $O(n)$ (which is also the complexity of $N(C)$). Operations $N(C_T)$, $N(C_S)$ and the test statistics T_C are all $O(1)$.

The resulting computational complexity of the algorithm is the pre-processing time, $O(n \log n)$ for time rank and $O(n^2 \log n)$ for space rank, plus the space-time cylinders generation which has complexity

$$|\mathbf{SL1}|(O(\log n) + |\mathbf{SL2}|(O(\log n) + |\mathbf{SL3}|[O(\mathbf{SS1}) + O(\mathbf{SS2}) + O(\mathbf{SS3})])).$$

where $|\mathbf{S}|$ denotes the number of times loop **S** is executed. From the previous discussion, $|\mathbf{L1}| = O(n)$, $|\mathbf{L2}| = O(n^2)$ and $|\mathbf{L3}|[O(\mathbf{S1}) + O(\mathbf{S2}) + O(\mathbf{S3})] = O(n)$. Therefore, the final complexity of our scan algorithm is $O(n^4)$.

Bibliography

Cliff AD, Ord JK (1981) Spatial processes: models and applications. Pion Limited: London.

Cox DR, Hinkley DV (1974) *Theoretical Statistics*. Chapman and Hall: London.

Diggle PJ, Chetwynd AG, Häggkvist R, Morris SE (1995). Second-order analysis of space-time clustering. *Statistical Methods in Medical Research*, 4, 124-136.

Duczmal LH, Assunção RM (2003) A simulated annealing strategy for the detection of arbitrary shaped spatial clusters. *Computational Statistics and Data Analysis*, in press.

Jacquez GM (1996) A k nearest neighbor test for space-time interaction. *Statistics in Medicine*, 15, 1935-1949.

Knox EG (1964) The detection of space-time interactions. *Applied Statistics*, 13, 25-29.

Kulldorff M (1997) A spatial scan statistic, *Communications in Statistics - Theory and Methods*, 26 (6): 1481-1496.

Kulldorff, M (2001) Prospective time periodic geographical disease surveillance using a scan statistic. *Journal of the Royal Statistical Society, Series A*, 164, 61-72.

Kulldorff M, Heffernan R, Hartmann J, Assunção RM, Mostashari F. A space-time permutation scan statistic for the early detection of disease outbreaks. Manuscript, 2003.

Kulldorff M, Hjalmar U (1999) The Knox method and other tests for space-time interaction *Biometrics*, 55 (2): 544-552.

Kulldorff M and IMS Inc. SaTSacan v3.0: Software for the spatial and space-time scan statistics. Available at [http://www.satscan.org], 2002.

Mantel N (1967) The detection of disease clustering and the generalized regression approach. *Cancer Research*, 27, 209-220.

Rogerson PA (1997) Surveillance systems for monitoring the development of spatial patterns. *Statistics in Medicine*, 16, 2081-2093.

Rogerson PA (2001) Monitoring point patterns for the development of space-time clusters. *Journal of the Royal Statistical Society, Series A*, 164, 87-96.