

Detecting Clusters in Spatially Repetitive Point Event Data Sets

La Détection de concentrations dans les semis d'événements ponctuels à répétitions spatiales

[Allan J. Brimicombe](#)

Centre for Geo-Information Studies, University of East London,
University Way, London E16 2RD, UK

Abstract : The analysis of point event patterns has a long tradition. Of particular interest are patterns of clustering or 'hot spots' and such cluster detection lies at the heart of spatial data mining. Certain classes of point event patterns have a significant proportion of the data having a tendency towards exact spatial repetitiveness. Examples are crime and traffic accidents. Spatial superimposition of point events challenges many existing approaches to cluster detection. In this paper a variable resolution approach, Geo-ProZones, is applied to residential burglary data exhibiting a high level of repeat victimisation. This is coupled with robust normalisation as a means of consistently defining and visualising the 'hot spots'.

Key-words : spatial clustering, point event data, spatial repetition, Geo-ProZone analysis, robust normalisation

Résumé : L'analyse des événements ponctuels a une longue tradition. La recherche de concentrations dans les semis d'évènements ponctuels ont un intérêt particulier, et la détection de ces concentrations est au cœur du data mining spatial. Certains modèles d'évènements ponctuels ont une proportion significative des données ayant une tendance vers la répétition spatiale exacte. Comme exemples on peu citer des crimes et des accidents de trafic. La superposition spatiale des événements ponctuels rend problématique beaucoup d'approches existantes pour détecter ces concentrations. Dans cet article une approche de résolution variable, des Geo-ProZones, est appliquée aux données de

cambrilage de residences montrant un niveau élevé de répétition spatiale. Ceci est couplé avec la normalisation robuste comme moyen de définir et de visualiser uniformément les concentrations.

Mots clés : analyse spatiale de proximité, événement ponctuel, répétition spatiale, analyses Geo-ProZone, normalisation robuste

1. Introduction

The analysis of point event patterns in geography, ecology and epidemiology has a long tradition (e.g. Snow, 1855; Clark & Evans, 1954; Harvey, 1966; Mantel, 1967; Cliff & Ord, 1981). The patterns detected are usually broadly classified as random, uniform or clustered. Although spatial randomness has traditionally been assumed to have no underlying process of interest, Phillips (1999) has nevertheless pointed out that such apparent randomness may be attributable to chaotic deterministic patterns and should therefore not be ignored out of hand. Where the point pattern shows spatial uniformity, a space-filling mutual exclusion process can be hypothesised. It is, however, clustered patterns that generally raise the strongest interest and hypotheses for underlying processes. Thus cluster detection lies at the heart of spatial data mining (Murray & Estivill-Castro, 1998; Openshaw, 1998; Murray, 2000; Miller & Han, 2001).

Clustered point patterns can be visualised as local concentrations of events in close proximity to one another, spatially separated by less dense or apparently random patterns of point events. However, certain classes of point event patterns have a significant proportion of the data having a tendency towards exact spatial repetitiveness (within the resolution of geo-positioning) although with a temporal separation between events. Typical examples include: crimes recorded against a property address (e.g. residential burglary, shoplifting, intimate partner violence), traffic accidents recorded against a section of road or intersection, utility failures recorded against a node or discrete section of network and so on. The focus in analysing such data sets is on defining 'hot spots' or 'black spots' where spatial clustering exists, but the occurrence of this spatial superimposition of point events challenges many existing approaches to cluster detection. In this paper a variable resolution approach, *Geo-ProZone analysis*, is applied to residential burglary data exhibiting a high level of repeat victimisation. This is coupled with *robust normalisation* as a means of consistently defining and visualising the 'hot spots'.

2. Cluster detection of 'hot spots'

The literature on clustering of point event data can be broadly classified into two approaches. One set of approaches is allied to mainstream statistics emanating from the work of Sokal & Sneath (1963). Thus clustering is a means of classification or grouping where clusters can be defined as “groups of highly similar entities” (Aldenderfer & Blashfield, 1984, p7). Spatially, this approach to cluster analysis will seek to form a segmentation into regions which minimise within-cluster variation but maximise between-cluster variation. There is a general expectation that the clustering will mutually exclusively include all points and is therefore space-filling within the geographical extent of the data (see for example Murray & Estivill-Castro (1998), Murray (2000)). Halls *et al.* (2001) and Estivill-Castro & Lee (2002) provide examples of the use of Dirichlet and Delaunay diagrams respectively to define spatial clusters. These algorithms, however, will fail where points occupy the same location. To delete duplicate points to overcome this problem will lead to important data loss, whilst to shift points slightly into non-duplicate positions will introduce bias away from being able to detect such repeat events. The other broad set of approaches uses spatially exhaustive search to identify localised excesses of event occurrences. Typical of this approach is the Geographical Analysis Machine (GAM) and its descendants (Openshaw *et al.*, 1987; Openshaw, 1998). Similar approaches are based around kernel density estimation (Silverman, 1986; Brunson, 1995) in which the highest densities form 'hot spots' (e.g. Gatrell *et al.*, 1996). This approach is particularly popular in crime analysis (Harries, 1999; Ratcliffe & McCullagh, 1999; McLafferty *et al.*, 2000) with GIS functionality available, for example, in the Spatial Analyst extension to ArcView® and in Hotspot Detective for MapInfo® (Figure 1).

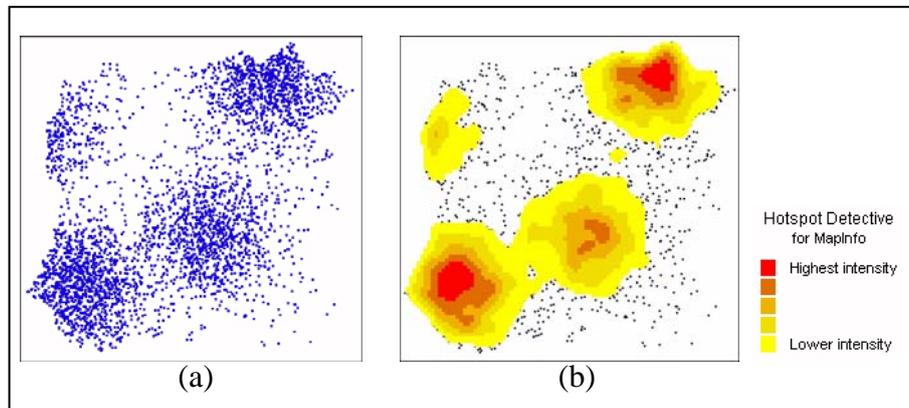


Figure 1. Kernel density estimation for 'hot spot' detection: (a) burglary point event data set; (b) kernel density estimation using default parameters (superimposed on point pattern) - 'hot spots' are usually taken to be the highest intensity locations.

The popularity of the kernel density estimation (KDE) approach is clear from its ease of use and the striking visualisations that it produces. It is nevertheless an interpolation that transforms the point events into a more-or-less smoothed continuous surface and with any such technique, parameters need to be set which are critical to the outcome. For KDE these are the underlying grid size and the kernel bandwidth. Reasonable values for parameters can be difficult to estimate and are often done so subjectively (Sabel *et al.*, 2000). Fotheringham *et al.* (2000) suggest an optimum bandwidth calculated from the standard distance. For situations where there are contrasting densities across a study area (e.g. urban to rural), an adaptive bandwidth can be employed (Brunsdon, 1995). Best practice would suggest a form of sensitivity analysis to identify optimum parameter values (Brimicombe, 2003). Figure 3 shows such an approach for a fixed grid size (one hectare) and varying bandwidth. The maximum nearest neighbour distance (NND) between point events in Figure 1(a) is 574m or approximately 12 times the median NND of 47.5m; so the bandwidth has been bracketed at three, six and nine times the median NND. The effect is to produce increased size and severity of 'hot spots'! What then is the truth? Although from a research perspective the sensitivity to grid size should also be tested, the pragmatics of the workplace usually means that analysts accept the default values for parameters suggested by the software as a matter of convenience.

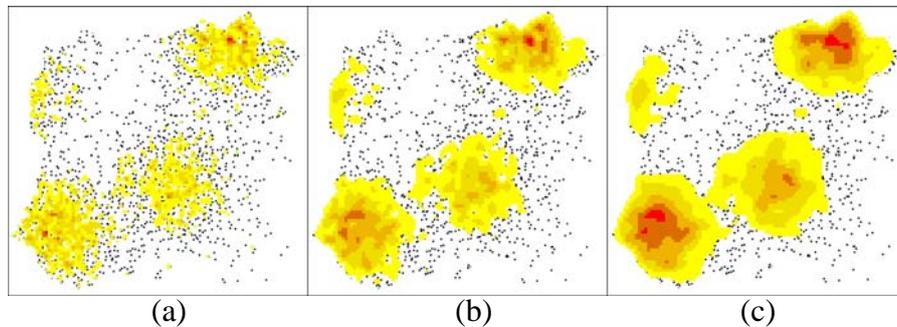


Figure 2. Searching for optimum bandwidth: (a) 3 times median NND; (b) 6 times median NND; (c) 9 times median NND; {for legend see Figure 1}.

The burglary data set presented here has a large number of repeat victimisations giving superimposed point events. Theory would suggest that 'hot spots' would be quite localised. High crime areas are primarily so because they are areas of high repeat offending and high repeat victimisation (Trickett *et al.*, 1992; Townsley *et al.*, 2003). KDE, as used by many police analysts, smoothes over the very localised repeat victimisations in favour of the regional pattern with choice of end result driven more by the aesthetic qualities of the visualisation. Boundary effects around the edge of data sets are also a problem for density estimation and perhaps not surprisingly police analysts tend not to find 'hot spots' at the edge of their jurisdictions. KDE software in the public domain by Atkinson & Unwin (2002) for MapInfo[®] does offer a buffer to avoid spurious values at boundaries but does not entirely overcome the problem of how to identify real 'hot spots' that exist at boundaries. Figures 1 and 2 focus attention on crime *counts*, that is, an elevated share of crime in a localised area. 'Hot spots' based on counts inform the deployment of resources in response to events. Less common in crime analysis (but more common, for example, in epidemiology) are 'hot spots' based on elevated *rates*. Such 'hot spots' reflect the level of risk and thus inform deployment of resources for mitigation. For the same distribution of point events, 'hot spots' based on counts are often different to those based on rates as the latter are not just a function of the distribution of point events but also of the underlying population at risk. Ideally both should be used. Sabel *et al.* (2000) report the use of KDE in association with an underlying population at risk to map

relative risk of disease occurrence. Whilst readily implemented, it suffers from the added difficulty of parameter estimation for two KDE surfaces (disease occurrence and population at risk) that are then combined to produce a ratio surface.

3. The Geo-ProZone algorithm

The theory of adaptive recursive tessellations is given in Tsui & Brimicombe (1997a) with applications of their use for spatial analysis in Tsui & Brimicombe (1997b). Specific application to point pattern analysis can be found in Brimicombe & Tsui (2000) and Brimicombe (2003). At the heart of adaptive recursive tessellations is a variable resolution approach to space. No longer are scale and resolution treated as being uniform across an area but are allowed to vary locally in response to the point pattern. This is achieved through a recursive decomposition of space, similar to quadtrees, but allowing variable decomposition ratios (quadtrees only have 1:4 ratio) and rectangular cells (quadtrees are usually restricted to square cells). The algorithm makes no prior assumptions about the statistical or spatial distribution of points. Each point is treated as a binary occurrence of some phenomenon without further descriptive attributes. The starting point is a convex hull of all the point events. Maximum x and y values of the data set form the maximal cell. The decision to further decompose any one cell larger than the atomic cell size is based on the variance at the next level of decomposition and a heuristic on the number of empty cells that result. The atomic cell size (or smallest possible cell size) is mediated between the median nearest neighbour distance and average cell size per point, whichever is smallest. Any cells formed through decomposition that fall outside the convex hull are automatically deleted. Tests have shown the algorithm to be consistently effective in comparison with other approaches of point cluster detection (see Brimicombe & Tsui 2000). The resulting clusters are termed Geo-ProZones (GPZ) as they represent zones of geographical proximity in the point pattern. As with kernel density estimation, the highest densities can be taken as 'hot spots'. However, GPZ are not an interpolation, but a segmentation into polygons having internal consistency in the distribution and density of the point events within them. Also, it does not suffer from boundary effects. GPZ for the burglary events in Figure 1 are given in Figure 3.

The pattern in Figure 3(b) reflects the pattern in Figure 1(b). The underlying speckle arises because all events are mapped without smoothing. The highest densities, or what would be interpreted as 'hot

spots', occur as more localised concentrations of repeat victimisation. Since GPZ results in polygons, they can be readily overlaid on an underlying population at risk (such as from census data) and reclassified as rates. Figure 3(c) shows GPZ as rates per thousand households from the underlying census data. The pattern of 'hot spots' is quite different and identifies where citizens are at greatest risk. Some of these areas appear reasonably extensive, others are quite localised where repeat victimisation is occurring.

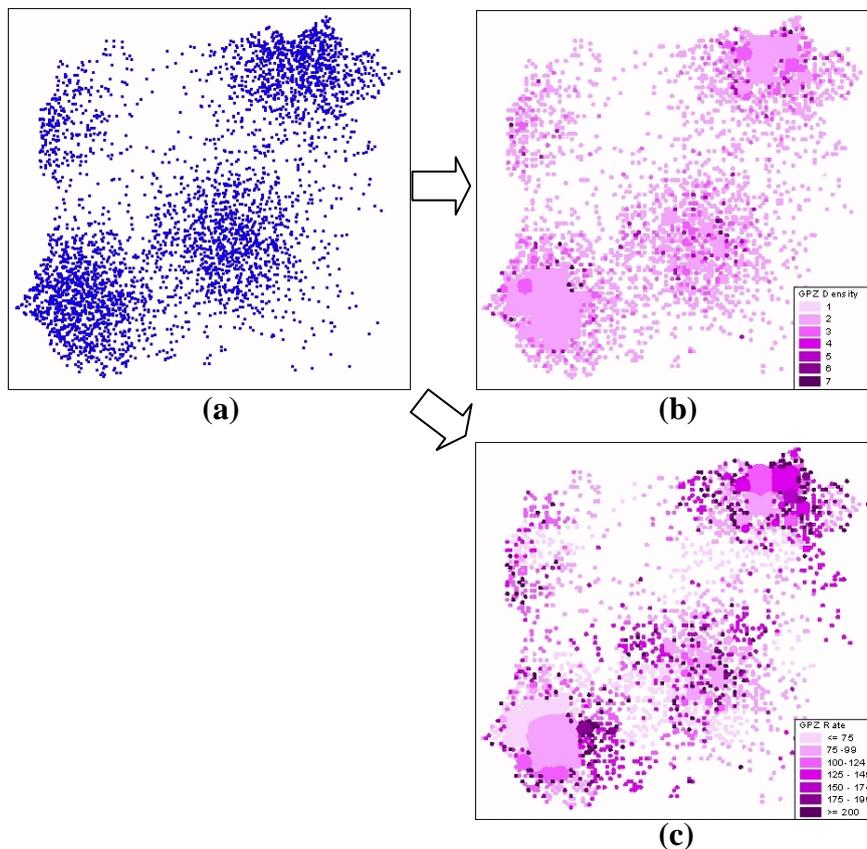


Figure 3. *Geo-ProZone analysis: (a) burglary point events; (b) GPZ for density of burglaries per hectare; (c) GPZ for rate of burglaries per thousand households.*

4. Robust normalisation for outlier detection and consistent visualisation

Whilst GPZ offers important methodological improvements in cluster detection where there is a tendency towards localised repetitive events, outstanding issues for this and any other approach relates to the well-known limitations of thematic mapping: number of class intervals, the fixing of class boundaries and what colours to use. There is the added issue of what constitutes the cut off for a ‘hot spot’. In practice, decisions often lack consistency. One approach is through data normalisation. A new form, robust normalisation [1], has been introduced (Brimicombe, 1999, 2000) as an alternative to the popular Z transformation where data are skewed.

$$RN = \frac{(x - median)}{(median - lower_quartile)} \quad \text{for } x < median$$

$$RN = \frac{(x - median)}{(upper_quartile - median)} \quad \text{for } x > median$$

$$RN = 0 \quad \text{for } x = median \quad [1]$$

The term ‘robust’ refers to the use of the median and inter-quartile range from robust statistics (Hettmansperger & McKean, 1998). The outcome of robust normalisation (Figure 4) is a distribution of median = 0, lower quartile = -1 and upper quartile = +1. Values of <-3 and >+3 are considered extreme values and the transformation can be used consistently for detecting outliers. It also provides a means of defining consistent class intervals and cartographic representation. Robust normalisation is achieved using the algorithm in [1] which is easily coded as a Microsoft® Excel macro.

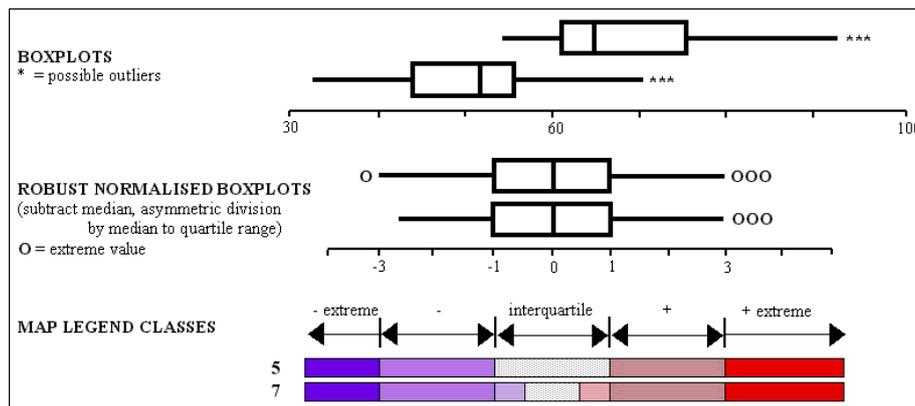
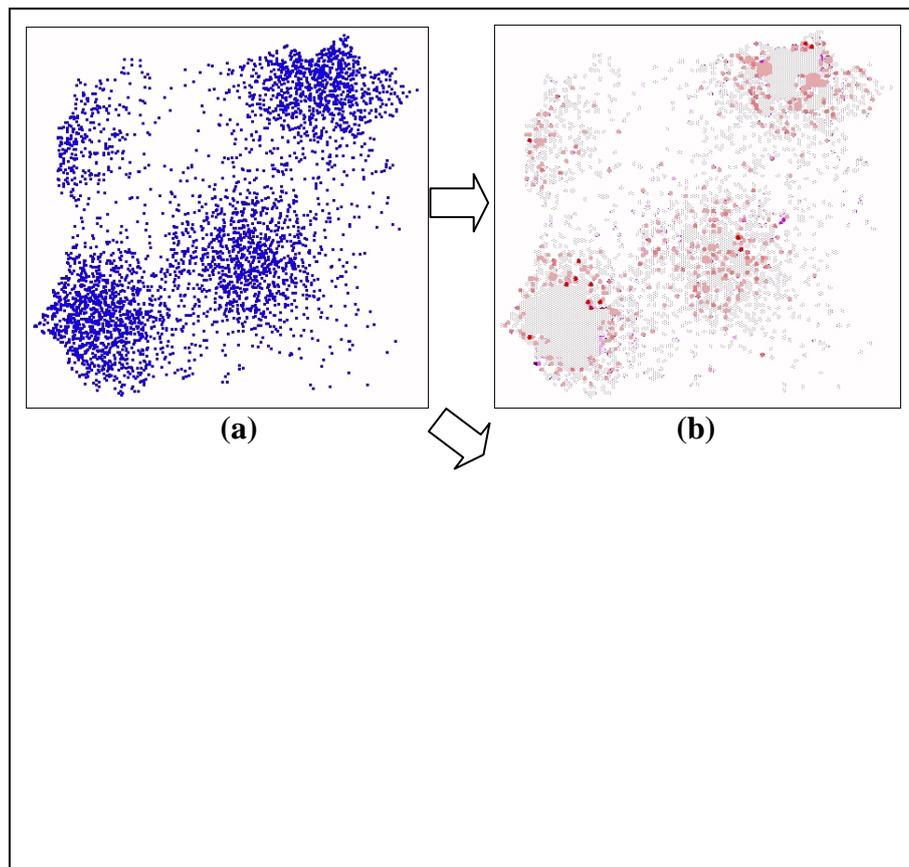


Figure 4. Robust normalisation of two dissimilar distributions to

achieve the same interquartile range, identify outliers and achieve consistent map legend classes (5 or 7 intervals) for visualisation.

Robust normalisation can be applied both to area-based data and to density estimate interpolations. For ‘hot spot’ detection it is the extreme positive values ($>+3$) that are of most interest. The robust normalised distribution easily lends itself to five or seven class intervals with class boundaries at quartiles (in the seven class interval scheme the values immediately around the median are further separated, as in Figures 4 and 5) and can be used in a standardised way for all visualisations. This allows more objective comparisons between maps. Figure 5 shows robust normalisation applied to both GPZ densities and rates from Figure 3 overcoming problems of arbitrary numbers of classes and class intervals. Localised ‘hot spots’ are where there are extreme positive values. By analogy ‘cool spots’ would be where there are extreme negative values. Clearly picked out in Figure 5(b) are the localised excesses of counts that represent ‘hot spots’ of repeat victimisation. More striking though is Figure 5(c) which shows many more localised excesses of rates when counts are related to the underlying population at risk. Importantly, in both cases the ‘hot spot’ distributions do not necessarily conform to initial subjective impressions of the point patterns (Figure 5(a)) as the ‘hot spots’ are in fact occurring where point events are superimposed and hence can neither be picked out by eye nor effectively by kernel density estimation.



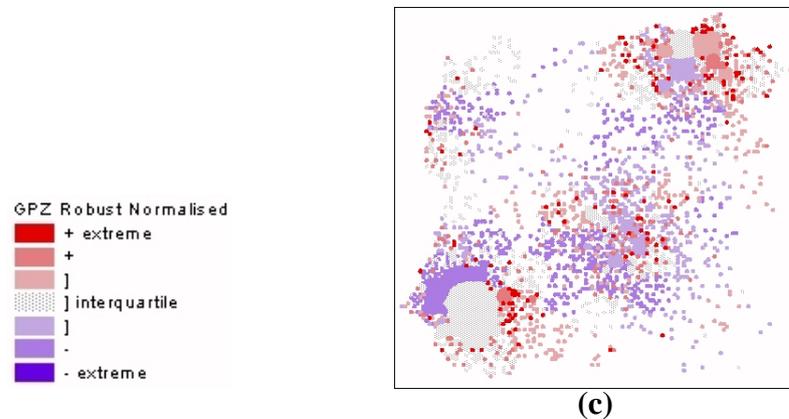


Figure 5. Applying robust normalisation: (a) burglary point events; (b) GPZ for density of burglaries per hectare; (c) GPZ for rate of burglaries per thousand households; {legend applies to both (b) and (c)}.

5. Conclusions

A consistent approach to cluster detection and visualisation of ‘hot spots’ through the combined use of Geo-ProZones and robust normalisation has been presented. The Geo-ProZones algorithm overcomes problems raised when data sets exhibit a tendency towards spatially repetitive events and where ‘hot spots’ will be highly localised. It also overcomes the boundary issues associated with KDE. The algorithm is suited to producing both segmentations of point densities and rates/risk in relation to underlying populations at risk. Problems can arise, however, from the presence of spatial outliers distorting the initial convex hull created around the point events. Improvements to the algorithm are being investigated to reduce sensitivity to any such outliers. Robust normalisation provides consistency in defining class intervals with ‘hot spots’ as localised extremes. Visual map comparisons for decision making are rendered more objective. Applications of the approach have been successfully used in analyses of crime, health and pipe bursts in water reticulation systems.

6. References

Aldenderfer, M. S., Blashfield, R. K., *Cluster Analysis*, California, Sage, 1984.

- Atkinson, P.J., Unwin, D.J., «Density and local attribute estimation of an infectious disease using MapInfo», *Computers and Geosciences*, vol. 28, 2002, p. 1095-1105.
- Brimicombe, A.J. «Small may be beautiful - but is simple sufficient?», *Geographical and Environmental Modelling*, vol. 3, 1999, p. 9-33.
- Brimicombe, A.J. «Constructing and evaluating contextual indices using GIS: a case of primary school performance», *Environment & Planning A*, vol. 32, 2000, p. 1909-1933.
- Brimicombe, A.J. *GIS, Environmental Modelling and Engineering*. London, Taylor & Francis, 2003.
- Brimicombe, A.J. «A variable resolution approach to cluster discovery in spatial data mining», dans Kumar *et al.* (eds) *Computational Science and its Applications*. Berlin, Springer-Verlag, vol. 3, 2003, p. 1-11.
- Brimicombe, A. J., Tsui H. Y. «A variable resolution, geocomputational approach to the analysis of point patterns», *Hydrological Processes*, vol. 14, 2000, p. 2143-2155.
- Brunsdon, C. «Estimating probability surfaces for geographical point data: an adaptive kernel algorithm», *Computers and Geosciences*, vol. 21, 1995, p. 877-894.
- Clark, P. J., Evans, F. C. «Distance to nearest neighbour as a measure of spatial relations in populations», *Ecology*, vol. 35, 1954, p. 445-453.
- Cliff, A. D., Ord, J. K. *Spatial Processes: Models and Applications*. London, Pion, 1981.
- Estivill-Castro, V., Lee, I. «Argument free clustering for large spatial point-data sets via boundary extraction from Delaunay Diagram», *Computers, Environment and Urban Systems*, vol. 26, 2002, p. 315-334.
- Fotheringham, S.A., Brunsdon, C., Charlton, M. *Quantitative Geography: Perspectives on Spatial Data Analysis*. London, Sage, 2000
- Gatrell, A. C., Bailey, T. C., Diggle, P. J., Rowlingson, B. S. «Spatial point pattern analysis and its application in geographical epidemiology», *Transactions of the Institute of British Geographers NS*, vol. 21, 1996, p. 256-274.
- Halls, P.J., Bulling, M., White, P. C. L., Garland, L., Harris S. «Dirichlet neighbours: revisiting Dirichlet tessellation for neighbourhood analysis», *Computers, Environment and Urban Systems*, vol. 25, 2001, p. 105-117.
- Harries, K. *Mapping Crime: Principle and Practice*. Washington DC, National Institute of Justice, 1999.
- Harvey, D. W. «Geographical processes and point patterns: testing models of diffusion by quadrat sampling», *Transactions of the Institute of British Geographers*, vol. 40, 1966, p. 81-95.
- Hettmansperger, T.P., McKean, J.W. *Robust Nonparametric Statistical Methods*. London, Arnold, 1998.
- MacQueen, J. «Some methods for classification and analysis of multivariate observations», *Proceedings of the 5th Berkeley Symposium on Maths and Statistics Problems* vol. 1, 1967, p. 281-297.

- Mantel, M. «The detection of disease clustering and a generalised regression approach», *Cancer Research*, vol. 27, 1967, p. 209-220.
- McLafferty, S., Williamson, D., McGuire, P.G. «Identifying crime hotspots using kernel smoothing», dans Goldsmith *et al.* (eds) *Analyzing Crime Patterns*, Thousand Oaks CA, Sage, 2000, p. 77-85.
- Miller, H. J., Han, J. *Geographic Data Mining and Knowledge Discovery*. London, Taylor & Francis, 2001.
- Murray, A. T. «Spatial characteristics and comparisons of interaction and median clustering models», *Geographical Analysis*, vol. 32, 2000, p. 1-18.
- Murray, A. T., Estivill-Castro, V. «Cluster discovery techniques for exploratory spatial data analysis», *International Journal of Geographical Information Science*, vol. 12, 1998, p. 431-443.
- Openshaw, S. «Building automated geographical analysis and explanation machines», dans Longley *et al.* (eds.) *Geocomputation: A Primer*. Chichester, Wiley, 1998, p. 95-115.
- Openshaw, S., Charlton, M. E., Wymer, C., Craft, A. W. «A mark I geographical analysis machine for the automated analysis of point data sets», *International Journal of Geographical Information Systems*, vol. 1, 1987, p. 359-377.
- Phillips, J. D. «Spatial analysis in physical geography and challenge of deterministic uncertainty», *Geographical Analysis*, vol. 31, 1999, p. 359-372.
- Ratcliffe, J.H., McCullagh, M.J. «Hotbeds of crime and the search for spatial accuracy», *Journal of Geographical Systems*, vol. 1, 1999, p. 385-398.
- Sabel, C.E., Gatrell, A.C., Löytönen, M., Maasilta, P., Jokelainen, M. «Modelling exposure opportunities: estimating relative risk for motor neurone disease in Finland» *Social Sciences & Medicine* vol. 50, 2000, p. 1121-1137.
- Silverman, B.W. *Density Estimation for Statistics and Data Analysis*. New York, Chapman & Hall, 1986.
- Snow, J. *On the Mode of Communication of Cholera*. London, Churchill Livingstone, 1855.
- Sokal, R., Sneath, P. *Principles of Numerical Taxonomy*. , San Francisco, Freeman 1963.
- Townsley, M., Homel, R. and Chaseling, J. «Infectious burglaries. A test of the near repeat hypothesis», *British Journal of Criminology*, vol. 43, 2003, p. 615-633.
- Trickett, A., Osborne, D.K., Seymour, J., Pease, K. «What is different about high crime areas?», *British Journal of Criminology*, vol. 32, 1992, p. 81-90.
- Tsui, H. Y., Brimicombe, A. J. «Adaptive recursive tessellations (ART) for Geographical Information Systems», *International Journal of Geographical Information Science*, vol. 11, 1997a, p. 247-263.
- Tsui, H. Y., Brimicombe, A. J. «Hierarchical tessellations model and its use in spatial analysis», *Transactions in GIS*, vol. 2, 1997b, p. 267-279.

Statement on trade names, trademarks: Where a name is used by a company to distinguish their product which they may claim as a trade name or trademark, then that name appears here

Cybergeo : Revue européenne de géographie, N° 387, 11 juillet 2007

13

with an initial capital or all in capital letters. Mention of such proprietary or commercial products does not constitute an endorsement, nor indeed a refutation of these products.

Allan Brimicombe is Professor and Head of the Centre for Geo-Information Studies at the University of East London, UK. Research interests include: spatial data mining, GIS and simulation modelling, location-based services. He is a Chartered Geographer.