INT. J. GEOGRAPHICAL INFORMATION SCIENCE, 1998, VOL. 12, NO. 7, 751–764

Research Article

Aoristic[†] crime analysis

JERRY H. RATCLIFFE
Geography Department, University of Nottingham, Nottingham, England, UK
email: jerry.ratcliffe@nottingham.ac.uk

and MICHAEL J. McCULLAGH
Geography Department, University of Nottingham, Nottingham, England, UK
email: michael.mccullagh@nottingham.ac.uk

(*Received 14 January 1998; accepted 17 May 1998*)

Abstract. Crime pattern analysis has tended to focus on 'hotspot' analysis techniques; the identification of areas with higher densities of criminal activity. This paper documents a different approach to determination of hotspots and aims to present a conceptual framework for the temporal analysis of aoristic crime data. This analysis monitors the change in crime patterns over time and can be applied to arbitrary areas, especially those smaller than police beat boundaries, which have been traditionally the smallest resolution areal unit studied within most British police forces. It also clarifies examination of lower crime areas, which can be overlooked in other forms of analysis. Crime data often lacks temporal definition and three different methods of temporal search technique are compared. Results from a new aoristic approach highlight a weekly Monday peak in motor vehicle crime on one division of Nottinghamshire Constabulary. In another example a historically weighted temporal model helps identify rising or falling crime rates.

1. Introduction

The encoding of police crime records has created opportunities for computerised crime pattern analysis where only variations on pin-mapping existed before (Hirschfield *et al.* 1995). The digital revolution has now been embraced by police forces and much more commercial and academic effort is currently being expended to find new methods of crime analysis. The spatial element of crime has also been recognised and geographical information is often linked with crime data records (Ratcliffe and McCullagh 1998) enabling GIS to become an additional tool in the war on crime. Work in spatial analysis has emphasised the detection of concentrations of crime incidents. STAC (Spatial and Temporal Analysis of Crime) is one system developed in Chicago which offers a simple system to detect crime 'Hot Clusters' (Illinois Criminal Justice Information Authority 1998). This type of analysis is more commonly referred to as 'hotspot' analysis. A map of crime locations is searched for

[†] Aoristic, without defined occurrence in time, *from*; Aorist (SOED); one of the past tenses of the Greek verb, which denotes a simple past occurrence, with none of the limitations of the other [past] tenses.

1365-8816/98 $12.00

areas of increased density, sometimes moderated against a background of population density or other factors. The theory shares similarities with studies into epidemiology and many of the tools originate in the medical field. Examples of this include Openshaw's GAM (Openshaw and Charlton, 1987) and the development of new spatial algorithms (Gatrell *et al.* 1996). With these tools, GIS has been used for crime mapping and analysis with varying degrees of success (Hirschfield *et al.* 1995, 1997, Openshaw *et al.* 1990). The temporal aspect of crime pattern analysis has received less attention even though active police officers know where crime is, but not necessarily whether it is rising or falling. The analysis presented here may allow them to monitor changing patterns obscured by other methods.

Temporal analysis of crime already takes place at national and regional levels. The annual publication of Home Office crime figures is a regular media event and local communities can compare the rise and fall in crime figures with the results for previous years at the county and national level. On a local level a crime analyst at a police division might plot the change in numbers of crime 'events' across the various beats of a station *if* the beat information is included with the crime record. The term 'event' has become accepted as a way of distinguishing the position of an individual observation within the study area (Gatrell *et al.* 1996). Temporal analysis is valuable for a number of reasons. There may be value in detecting the change in crime level over time and moving the focus away from those areas with the highest concentrations in crime. Lower concentration areas may be interesting if the crime rate has suddenly risen or fallen, not in relation to other areas, but in relation to the historical pattern in that area. If changes are evident, then one explanation may be presented by Cohen and Felson's routine activity theory (Cohen and Felson 1979) which argues that most crime requires the convergence in space and time of a motivated offender, a suitable target and the absence of a capable guardian. Spatially constrained non-temporal hotspot analysis can obscure the variation in areas with lower levels of crime. Until the introduction of detailed spatially referenced crime data the police beat was generally the smallest resolution available for temporal and spatial analysis. This article aims to suggest a methodology for variable scale temporal analysis of crime data and in the process identifies some of the benefits and pitfalls which the user can expect from applying aoristic spatiotemporal process models. Examples from a police division of Nottinghamshire Constabulary provide real world case studies of the process in action. During the course of the study a significant temporal relationship was uncovered, unobserved using other analytical techniques.

2. Temporal analysis

Increasingly it has been recognised that data models should incorporate a time element to reflect the changing nature of the world (Raafat *et al.* 1994). Much of the research considered varying methodologies of recording change within a temporal database and the problems of how much historical information to retain and in what format (Langran 1989, 1993, Peuquet and Niu 1995). These papers contain some of the more thorough reviews of recent temporal analysis. In one article, Donna Peuquet presents a model for representing temporal change and likens the passage of time to an elastic line, with events appearing as knots at different places along the time-line (Peuquet 1994). This is a useful analogy and allows us to conceptualise temporal events as having a start point, a duration and an end point. Temporal topological relationships can be defined with Boolean operators used to quantify the dynamics of any interaction between events (Peuquet and Niu 1995). Different organisational approaches include time based methods where the temporal dimension is the main organisational factor, and the more straight-forward 'snapshot' approach presented here.

The snapshot method is not only the most easily comprehended data model, it is also fairly easy to integrate into current GIS. A sequence of spatially-registered grids is employed to record the 'world state' at a given time. At different times the same grid is employed to record the 'world state' at another time and in this manner a collection of images is recorded, each at a unique time. Figure 1 shows S_i, a snapshot of a given state taken at a time t_i. As Peuquet and Niu point out, the temporal distance between snapshots does not necessarily have to be uniform. Each pixel on the grid records the state of the desired variable at the chosen time, for just that location within the grid. It should be noted that each 'snapshot' records a given state and not the changes that have taken place. These methods are most effective when temporal events have a definite structure. This is unfortunately rarely the case with crime incidents.

3. Temporal analysis of crime data

Temporal queries of crime databases are more involved because of the complex nature of crime recording. Unless the criminal is disturbed or captured, it is unlikely that the exact time of the offence will be known. Although rarely an issue for crimes like robbery (where victims tend to know when they were robbed) this is a problem for incidents like motor vehicle theft and burglary—crimes which are often identified as high prevention and detection priorities for police forces. Daytime burglaries can occur when the occupants are at work and are episodes which rarely take longer than a few minutes. When interviewed by the police the victims are unable to narrow the incident down to any time more specific than a number of hours. Police crime recording practices reflect this by documenting a number of fields for the time of incident. Although field names vary from force to force, variations on *on_date*, *at_time*, *from_date*, *from_time*, *to_date* and *to_time* allow the crime record sheet to incorporate a range of possible incident times. This can vary from an exact event time to a number of weeks or longer, if the unfortunate occupants of a burgled premises have been on holiday, for example. This creates problems when it is desired to search a crime database for events that occurred during a specific time period. A number of solutions to this problem are possible, however each has limitations on its functionality. During the course of this study three different methods of temporal search were considered; an averaging temporal search which averages the date and time fields, a rigid temporal search which only contains definite records within the

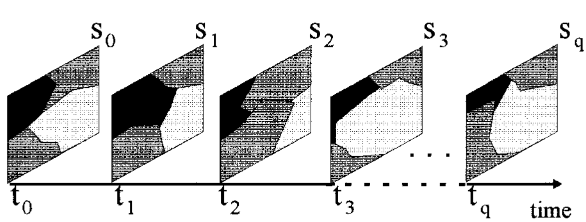

Figure 1. Snapshot temporal analysis approach (Peuquet and Niu 1995, with permission).

search criteria, and an aoristic search which considers all records which *might* have occurred within the search criteria time.

One of the simplest solutions is to take the average of the start and end event times. For example, if the victims of a burglary leave their house at noon and return to their burgled house at 6pm, the time of 3pm would be used by an averaging temporal method for any crime analysis. This process is summarised in figure 2 where horizontal bars represent crime incidents which have a start time, a duration (the length of the bar) and an end time. The search criteria similarly has a start and end time shown by the vertical arrows. The small vertical arrows in the middle of each event represent the location of the average along the time line. Although a considerable amount of the second incident takes place within the search parameters, it is not included because the location of the average is outside the search criteria. Averaging the date field has been used in studies of the burglary repeat victimisation (Johnson *et al.* 1997). Although computationally simple and an adequate solution for the longer time periods of repeat victimisation studies, this is a compromise answer and ties the event time to one possible time which, although the mean of the possible event times, is no likelier than any other time. A search based on fixed temporal windows adds emphasis to only one search period and denies the incident the chance to register with other equally applicable categories. The advantage of this system is that each event will be recorded as one definite time and the results of the temporal query should have the same number of incidents as the originating database.

Another possibility is to record those incidents which fall within the search period. This will be called the rigid temporal search method. Search queries have to be constructed more carefully to reflect the more inflexible nature of the search engine. This process is summarised in figure 2 where the horizontal bars again represent crime incidents which have a start time, a duration (the length of the bar) and an end time. The search criterion similarly has a start and end time shown by the vertical arrows. Only the fourth incident is returned by the search method. All of the others have timelines which extend beyond the confines of the search criteria. The outcome of this type of search method is a result with generally lower numbers than the originating database, but with a higher degree of accuracy in the temporal search.

Aoristic crime presents a particular problem in that it is temporally uncertain and can fall either side of the search boundary. The third approach, called here the aoristic search method, is to record crime incidents that *might* have occurred within the search time. Conceptually it is more complicated because a single incident can simultaneously register in a number of search categories if it covers an extended

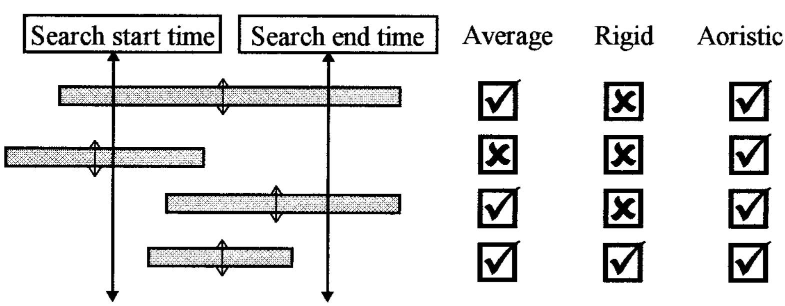

Figure 2. Summary of averaging, rigid and aoristic temporal search processes.

time period. In this manner an incident will appear in each time category where the slightest possibility exists that it might have occurred in that search block, however small a percentage of the incident time covered the search category. The resulting table will show a larger number of 'hits' than existed in the originating database, each 'hit' registering the *possibility* of an incident rather than a definite incident as in the previous method. This process is also summarised in figure 2.

The second and third approaches offer more accurate and elegant solutions than the averaging method first described. The averaging temporal search method is acceptable for long term studies such as repeat victimisation and annual or monthly comparisons, but for shorter time scales it lacks accuracy and can easily generate erroneous results because it fixes the crime incidents at a totally arbitrary date and time. The second more rigid method presents the more precise means of identifying temporal events, while the last method allows for the exploration of all possible events from the originating database.

As their name suggests, the rigid and aoristic search methods offer different interpretations of the data in their results. The following factors should be considered in choosing which method to use. The size of the database generated might have repercussions for later statistical operations. The temporal vagueness of the original database is also important. The more accurate the temporal fields in the database, the less a method like aoristic search becomes necessary.

Figure 3 and figure 4 were generated from a daily analysis of crime records from Trent division, Nottingham Constabulary in November 1996. Trent division is a large division of Nottinghamshire Constabulary situated in the South and West of the City of Nottingham, the beats of which can be seen in the right-hand diagram on figure 6. Both aoristic and rigid methods were employed to examine the daily crime rate for assault and motor vehicle crime. For definition purposes assault includes all counts of actual and grievous bodily harm. Car crime includes all counts of motor vehicle theft, theft from motor vehicle, and TWOCing (taking without consideration)—often referred to as joyriding.

Figure 3 shows the number of records recovered by both the aoristic and rigid search methods when querying all assaults on Trent division during November 1996.

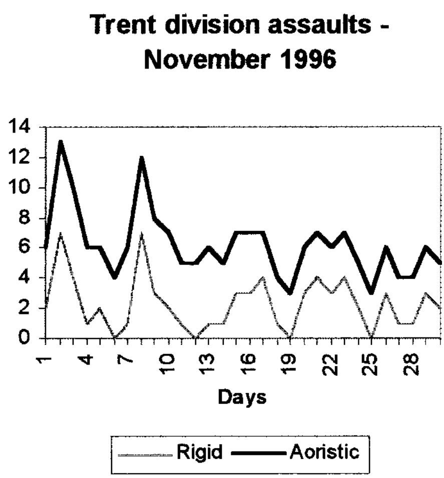

Figure 3. Rigid and aoristic search of assaults over 1996.

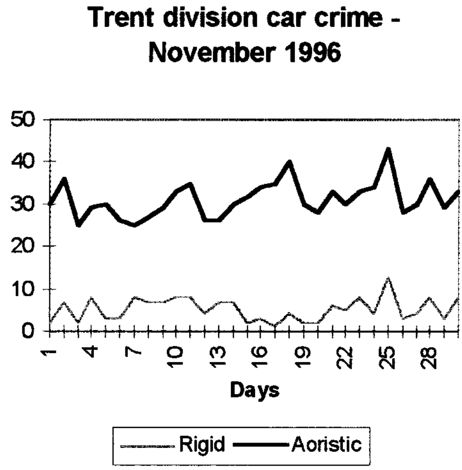

Figure 4. Rigid and aoristic search of car crime over 1996.

The curves show a strong positive correlation. The main difference between the results are the existence of three crime records which detailed cases of domestic assault lasting a period of years. In these circumstances the police often use one crime sheet to record an ongoing catalogue of incidents. Removal of these three incidents from the graph would bring the aoristic and rigid crimes into almost exact correspondence. With this type of crime it is clear that the rigid search method is sufficient. Figure 4 shows the difference between rigid and aoristic search processes for car crime over the same study area and time. Note the difference in vertical scale from figure 3. The absolute difference between the values is indicative of the lack of temporal accuracy in police crime reports due to the inability to tie down the time of incident. This temporal inaccuracy is most evident with the apparent regular 7-day peaks in the aoristic data. These occur on the 11, 18 and 25 of November 1996 and are not detectable in the rigid data (with the exception of the 25 November). The original data was re-examined in an attempt to discover why these Monday peaks existed. It was originally thought that the victims might have been away for the weekend and returned in the early hours of Monday morning to find their car stolen or stolen from. This was not the case and it appears from detailed crime record inspection that Sunday and Monday nights were simply the most popular with local car thieves in November 1996. This trend is not visible in the rigid temporal search method. A full statistical analysis was conducted with a larger time sample of data to determine if the Monday pattern was significant.

Six months' worth of car crime events were extracted from the crime database for both rigid and aoristic search methods and are displayed in figure 5. To prevent replication in the data, the results from the aoristic temporal search were weighted according to their temporal length to form the adjusted aoristic set shown in figure 5. A crime taking place on a single particular day contributes a score of one to the result, but an event with an end date four days after the start date would contribute 0.25 to each of the intervening days, as the exact day is unknown and there is no hypothesis as to the likely distribution of the occurrence of the event within those four days.

The selection of appropriate statistical analysis was complicated by the unknown

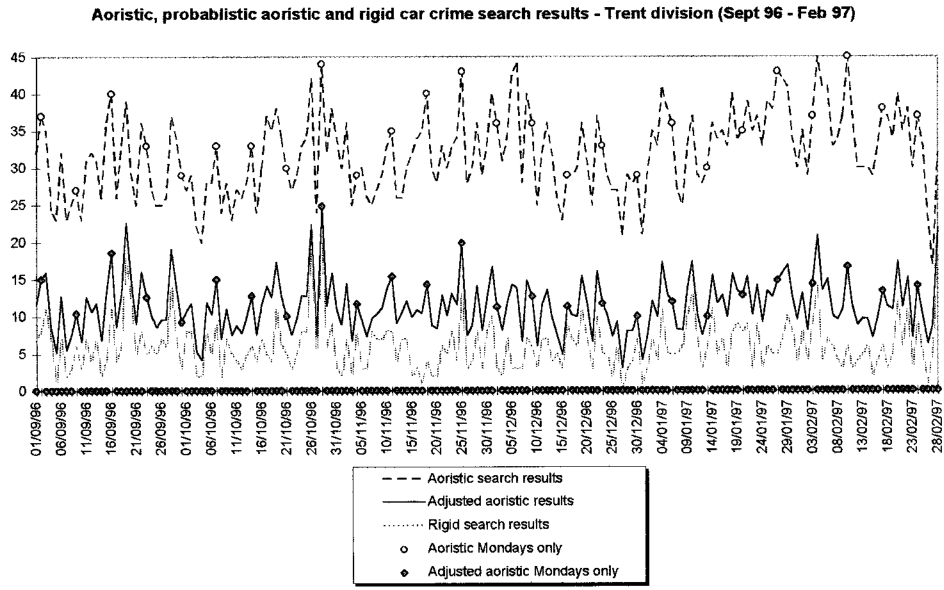

Figure 5. Daily count of car crime for Trent division for 6 month period.

distribution of the data, the potential lack of sample independence in the aoristic data sets, and uncertainty as to the underlying data generation process. Two tests were used; a standard single factor exploratory ANOVA, considered by some authors (e.g. Alder and Roessler 1962, p. 245) robust under adverse distributional conditions; and the less powerful non-parametric Kruskal-Wallis test which is designed to use non-normally distributed data of diverse variability. The problem of sample independence remains, but in the authors' view is not greatly infringed. A comparison of the results of the two approaches is given below.

Four ANOVA single factor (by day of week) tests were calculated, two using the rigid data and two using the adjusted aoristic data. Table 1 shows the number of car crimes committed each day of the week and the means and variances for both the rigid and adjusted aoristic data sets. Inspection of the table would suggest that more car crimes occur on Mondays. The null hypothesis was that Mondays were no different to other days of the week in terms of crime occurrence. The rejection level was set at 0.05. The results can be seen in table 2 for all days of the week for rigid and aoristic searches, and on all days of the week except Mondays for both

Table 1. Basic counts by day of week using adjusted aoristic search for Trent division car crime records.

Groups	Count	Sum	Mean	Variance
Monday	26	355.87	13.69	12.06
Tuesday	26	268.35	10.32	14.03
Wednesday	26	272.80	10.49	6.83
Thursday	26	316.77	12.18	13.42
Friday	26	281.92	10.84	13.27
Saturday	25	292.03	11.68	19.69
Sunday	26	288.60	11.10	10.63

Table 2. Single factor ANOVA tests of search results from Trent division car crime events—September 1996 to February 1997.

Test number	Test of variance between groups	df	F value (calculated)	F_{crit} value ($p=0.05$)	Probability level	Ho rejected
(1)	Rigid all days	6	1.94	2.15	0.08	No
(2)	Rigid not Mondays	5	1.45	2.27	0.21	No
(3)	Aoristic all days	6	2.80	2.15	0.01	YES
(4)	Aoristic not Mondays	5	1.02	2.27	0.41	No

rigid and aoristic search methods. Test results in table 2 were calculated for all days of the week for rigid (1) and aoristic (3) searches, and a further test was conducted on all days of the week except Mondays—again for both rigid (2) and aoristic (4) search methods. The purpose of (1) and (3) was to test whether significant Monday peaks occurred in the data, and in (2) and (4) to show that the other days of the week were undifferentiated.

There is no significant difference in either of the rigid search results (1) and (2). No variation in the pattern of crime, and specifically no Monday peak, has been detected by the rigid search method. Many records are excluded from the rigid search as they have timelines which extend beyond the daily search criteria (i.e. overnight). The aoristic method shows that Monday peaks are present (3). If Mondays are excluded (4), there is no significant difference between the daily means in the dataset and no pattern emerges. This study shows that there is a statistically strong Monday peak over the six month period. This result would not have come to light if only a purely rigid search for exact records had been performed. The results of the equivalent Kruskal-Wallis test are shown in table 3 for the aoristic data set alone. The same conclusions are reached with much the same significance levels as before. This would indicate some confidence in the results.

4. Effects of boundaries

This type of analysis has to use a geographical framework and small scale boundaries such as enumeration districts or generated grids are adequate for the task. Grid patterns are easy to generate in programmable GIS. Software houses such as MapInfo even provide freeware packages for grid creation and other functions which are available to download from their web site (http://www.mapinfo.com). Regular grids have an additional advantage in crime pattern analysis as they allow the generalisation of point data which preserves the anonymity of the original event locations. In this manner, results of crime analysis can be passed to outside agencies such as council crime prevention panels without compromising the sensitivity of the original data. Whatever the chosen boundaries, it is not a requirement that they be

Table 3. Single factor Kruskal-Wallis tests of search results from Trent division car crime events—September 1996 to February 1997.

Test number	Test of variance between groups	df	H value (Distributed as (Chi Square)	Chi square ($p=0.05$)	Probability level	Ho rejected
(1)	Aoristic all days	6	17.69	12.59	<0.01	YES
(2)	Aoristic not Mondays	5	5.76	11.07	>0.40	No

the same shape or volume. A temporal analysis will compare crimes within one polygon with the same boundary at another time, and not with neighbouring locations. For example, if two grid squares are next to each other, one of them could be mostly parkland. A search of motor vehicle crime in a parkland square will usually exhibit a low score. This is not an issue unless environmental change leads to the building of new housing estates where the parkland used to be. Grids can cause problems when the grid referencing system of the crime data is considered.

Within Nottinghamshire Constabulary most grid references are attached to crime records by querying the postcode and address and then attaching the AddressPoint 0.1 m resolution Ordnance Survey grid reference from a central gazetteer. However the AddressPoint catalogue at Nottinghamshire Constabulary is not complete and any unknown locations have the Postcode Address File (PAF) 100 m resolution grid reference merged with the crime event record as a back-up procedure. If a grid uses a similar 100 m resolution, a number of crime events can find themselves on the exact location of grid intersections and either included in two or more grid results, or excluded from any grid location dependent on the GIS areal search procedure employed. Using grid co-ordinates that are not rounded to the nearest 100 m reduces this effect, although complete alleviation of the problem is only possible by reprogramming the areal selection tool in the GIS used.

Figure 6 shows the study area, a one kilometre resolution grid square over part of West Bridgford, a suburb of Nottingham. A year's worth of data (1996) was used to produce a pattern of all crimes over 52 weeks. The grid corners were deliberately

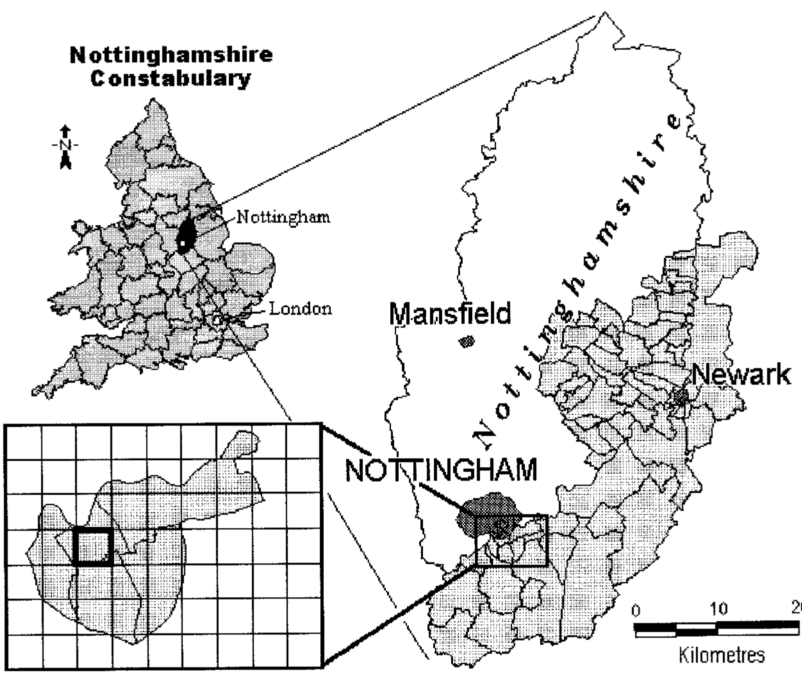

Figure 6. One kilometre resolution grid over West Bridgford station area, Trent Divison, Nottingham Constabulary, UK. Study cell within the grid is outlined. Large map shows the current Trent division beats—the smallest resolution for local spatial crime analysis prior to the use of GIS.

programmed to avoid the 100 m resolution problem mentioned earlier. To examine the extent of this problem, the grid was converted to polylines and the database interrogated to find how many crime events intersected the lines of the grid. From a database of 8,267 records, 62 were found to exactly intersect the lines of the grid, only 0.8%. With a sufficiently large database, this is an acceptably low percentage, but caution must be exercised if much smaller grids with finer resolution are employed, increasing the number of lines and raising the probability of intersection.

5. Analysis of mixed spatial and temporal studies

Useful off-the-shelf spatial and temporal crime analysis software is rare. Manual programming of a GIS with a Standard Query Language (SQL) is easily possible, though it can be laborious and time-consuming. Fortunately most proprietary GIS packages such as ArcView and MapInfo have programming languages which can automate many tasks. Once the temporal analysis process has been programmed, there are a variety of ways to analyse the results of a temporal crime pattern query.

Two realistic aims are to monitor the trend in crime rate for an area over time (often achieved using some form of regression analysis), and to identify time periods of particularly high or low rates to enable further investigation of their causes. The data takes on added complexity because the grid or boundary framework has a two-dimensional structure with time added as the third dimension. During this study it was felt that the most effective means of analysing the data was by using the variety of techniques available within a database program. Each geographical unit can be examined individually and the change in crime rate plotted against time.

Figure 7 shows the result of a year long sample of the target grid square when it was 'aoristically' searched for all crimes which occurred on a weekly basis during 1996. While regression analysis has applications in policework for mapping the general trend of the crime patterns in grid cells, one possible alternative could be an adaptive running mean. Figure 7 shows the count of all crimes in the target grid

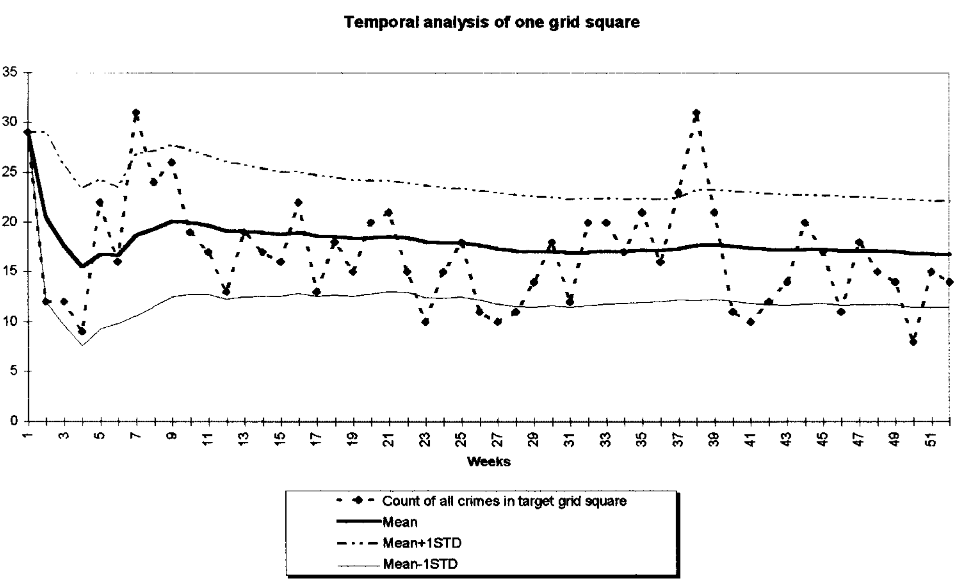

Figure 7. A simple adaptive monitoring model.

square for the same study period, as before. It also shows the running mean adjustment as each new value is received by the system. With every new value a revised standard deviation is calculated and two guidelines show the average plus and minus one standard deviation. This adaptive process enables monitoring of crime rates to determine highs and lows on a continuously variable basis.

The advantage of this system is that a threshold can be set and an alert given if exceeded. These unusual values can then be examined to find the cause. A 'settling down' period can be seen at the start during which time the system is of limited value owing to insufficient values being included in the mean. A further extension of this process was considered in this study by using a linear weighting system over the most recent 15 values in the running mean, weighted in a descending linear pattern which gives emphasis to the most recent values. The result can be seen in figure 8. There are obvious benefits to this system. The process is much better able to react to unusual values which are immediately absorbed by the system but lose significance the further they proceed away from the newest values down the time line.

After the initial 'settling-down' period, the unusual values at the start mean that the model has a high standard deviation, represented by the two lines either side of the central mean which show the mean plus and minus one standard deviation. Once more stable values are entered to the system for weeks 11 to 35, the standard deviation rapidly reduces and the lines converge on the mean, indicating that the unusual values are exerting less influence. Weeks 37 to 39 are again unusually high and the system adapts to reflect this, though the standard deviation is dropping by the end of the graph as more stable values are received by the system. Although a standard deviation is used here as the threshold, any measure of weighted sample variability could be employed.

6. Display of results

The results of the analysis of one grid cell were presented in figure 8, but the one kilometre resolution search area in figure 6 contained a 9 by 7 grid of 63 squares.

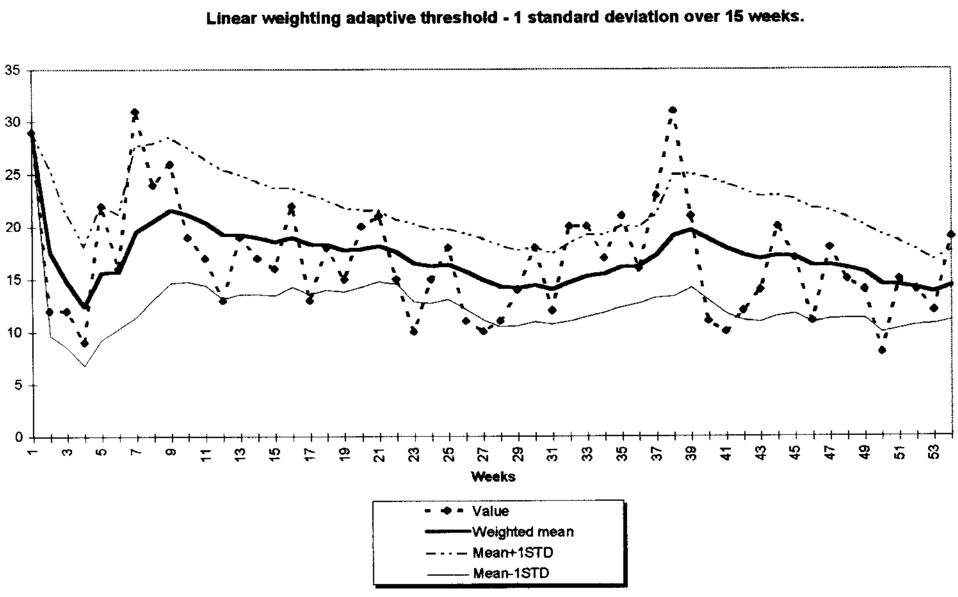

Figure 8. Linear weighting adaptive threshold process.

This presents a problem for presenting the results. Animation has been suggested as the solution to many GIS problems involving time as a variable (Dorling and Openshaw 1992, MacEachren 1994). This is a field where animation might be particularly useful. GIS packages can be easily programmed to output selected screens as image files of screen captures. A repeated number of these images can be moulded together into a FLI file, an AVI file or a number of other computer movie/animation formats. A number of programs are available to perform this operation, including a number of free or shareware products like Dave's Targa Animator (available through HENSA). It is beyond the scope of this paper to describe the full animation process, but a suggested animation sequence might be the definition of a map surface with the grid hidden over the top of it. As the sequence (and therefore the weeks) progresses grid cells only become 'active' if they are either above or below the tolerance threshold. Effective use of colour could code cells as one colour for being above and a different colour for being below the threshold. For a more thorough treatment of the complete animation process (see McCullagh 1998). This technique has a number of advantages. It is an effective means of displaying a large quantity of numerical information in an easily understandable format. It is also an efficient means of explaining the temporal aspect, and utilises a grid base which preserves the anonymity of the original data. Modern processors and algorithms have speeded up the animation process to the extent that a 100 frame animation can be created on a domestic PC in a few minutes.

The temporal change can be summarised in a more limited but accessible means by showing the number of threshold crossings within each grid square on a simple choropleth map. Figure 9 shows a number of choropleth maps of West Bridgford station beats, a sub-division of Trent division. Figure 9(*a*) shows the summary of all crimes on the sub-division for 1996. This type of map is one of the standard range of simple distribution maps commonly employed in police work, but which can mislead the user as to changes in more complex crime patterns. Figure 9(*b*) and figure 9(*c*) depict the number of times over the year that the grid squares aoristic temporal count exceeded the weighted running mean by more than one standard deviation. The combination of these two images in figure 9(*d*) shows areas with complex rapidly changing weekly crime rates, and also the areas with more stable crime event patterns. The temporal analysis in figure 9 allows a much more detailed interpretation of the temporal crime pattern and shows that changes in pattern are occurring in very different locations from those of major crime. The picture is much more complex than that given by the single display of total crime in figure 9(*a*).

7. Conclusions

This paper has presented a basic conceptual framework for aoristic analysis of crime data, and documented the process from accurate selection of the records through to presentation of some analytical results. Crime data frequently lacks temporal definition and two methods of performing temporal searches have been tested. Rigid temporal searches pick only those records which fall exactly within the search period, while aoristic methods can be used to explore a fuller range of temporal possibilities. The example of Monday car crime peaks brings to light patterns obscured by simpler methods, but begs the unresolved question of the causes of Monday car crime. The approach in this paper has been to examine the historical pattern and identify thresholds around present variation outside which abnormalities can be evaluated.

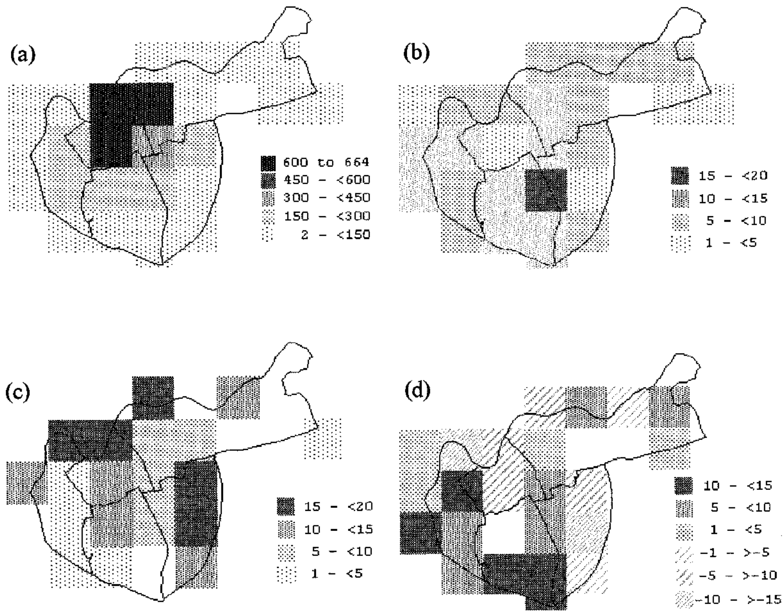

Figure 9. West Bridgford station beats overlaid with one kilometre square grids showing general crime and temporal analysis summaries. (*a*) Standard police crime mapping practice using totals for the year. (*b*) Aoristic search generated data set used to determine weighted running mean thresholded results. Here the number of times the upper threshold has been exceeded during the year is mapped. (*c*) The same data as for (*b*) but a count of the number of times the lower threshold has been crossed. (*d*) The difference between (*b*) and (*c*) to show the predominance of upper or lower threshold crossings. A negative class on the map indicates more significant reductions in crime and a positive class indicates significant increases.

Aoristic crime analysis is a method of charting the historical pattern over variable small areal units, independent of influence from neighbouring cells. Each cell can be plotted to examine the rise and fall of crime compared with a variable threshold over time. Temporal trends *in extrema* can be examined for patterns and can be correlated against other intelligence gathered at the police station. The adjusted aoristic search procedure, added to adaptive thresholding techniques, provides new tools for crime analysis.

Acknowledgment

The authors would like to thank the Nottinghamshire Constabulary for access to the data and co-operation in this research. The authors would also like to express their appreciation to the reviewers of this manuscript for helpful comments.

References

Alder, H. L., and Roessler, E. B., 1962, *Introduction to probability and statistics* (San Francisco, London: Freeman).

Cohen, L. E., and Felson, M., 1979, Social change and crime rate trends: A Routine Activity Approach. *American Sociological Review*, **44**, 588–608.

Dorling, D., and Openshaw, S., 1992, Using computer animation to visualize space-time patterns. *Environment and Planning B—Planning & Design*, **19**, 639–650.

Gatrell, A. C., Bailey, T. C., Diggle, P. J., and Rowlingson, B. S., 1996, Spatial point pattern analysis and its application in geographical epidemiology. *Transactions, Institute of British Geographers*, **NS 21**, 256–274.

Hirschfield, A., Brown, P., and Todd, P., 1995, GIS and the analysis of spatially-referenced crime data: Experiences in Merseyside, UK. *International Journal of Geographical Information Systems*, **9**, 191–210.

Hirschfield, A., Yarwood, D., and Bowers, K., 1997, Crime pattern analysis, spatial targeting and GIS: The development of new approaches for use in evaluating community safety initiatives. In *Crime and health data analysis using GIS* (Sheffield: SCGISA).

Illinois Criminal Justice Information Authority, 1998, *STAC user manuals*. (Chicago).

Johnson, S. D., Bowers, K., and Hirschfield, A., 1997, New insights into the spatial and temporal distribution of repeat victimization. *British Journal of Criminology*, **37**, 224–241.

Langran, G., 1989, A review of temporal database research and its use in GIS applications. *International Journal of Geographical Information Systems*, **3**, 215–232.

Langran, G., 1993, Issues of implementing a spatiotemporal system. *International Journal of Geographical Information Systems*, **7**, 305–314.

MacEachren, A., 1994, Time as a cartographic variable. In H. Hearnshaw and D. Unwin, eds., *Visualisation in Geographical Information Systems* (London: John Wiley), pp. 115–130.

McCullagh, M. J., 1998, Quality, use and visualisation in terrain modelling. In *Landform monitoring, modelling and analysis*, edited by S. N. Lane, K. S. Richards and J. H. Chandler (London: Wiley), Chapter 5, 95–117.

Openshaw, S., and Charlton, M., 1987, A Mark 1 Geographical Analysis Machine for the automated analysis of point data sets. *International Journal of Geographical Information Systems*, **1**, 335–358.

Openshaw, S., Cross, A., Charlton, M., and Brunsdon, C., 1990, Lessons learnt from a Post Mortem of a failed GIS. In *2nd National Conference and Exhibition of the AGI* (Brighton: AGI), 2.3.1–2.3.5.

Peuquet, D. J., 1994, It's about time—a conceptual framework for the representation of temporal dynamics in Geographical Information Systems. *Annals of the Association of American Geographers*, **84**, 441–461.

Peuquet, D. J., and Niu, D. A., 1995, An event-based spatiotemporal data model (ESTDM) for temporal analysis of geographical data. *International Journal of Geographical Information Systems*, **9**, 7–24.

Raafat, H., Yang, Z. S., and Gauthier, D., 1994, Relational spatial topologies for historical geographical information. *International Journal of Geographical Information Systems*, **8**, 163–173.

Ratcliffe, J. H., and McCullagh, M. J., 1998, Identifying repeat victimisation with GIS. *British Journal of Criminology*, **38**(4), 651–662.